

THE POWER OF PRIORITY CHANNEL SYSTEMS *

CHRISTOPH HAASE, SYLVAIN SCHMITZ, AND PHILIPPE SCHNOEBELEN

LSV, ENS Cachan & CNRS & INRIA, France
e-mail address: {haase,schmitz,phs}@lsv.ens-cachan.fr

ABSTRACT. We introduce Priority Channel Systems, a new natural class of channel systems where messages carry a numeric priority and where higher-priority messages can supersede lower-priority messages preceding them in the fifo communication buffers. The decidability of safety and inevitability properties is shown via the introduction of a *priority embedding*, a well-quasi-ordering that has not previously been used in well-structured systems. We then show how Priority Channel Systems can compute fast-growing functions and prove that the aforementioned verification problems are $\mathbf{F}_{\varepsilon_0}$ -complete.

1. INTRODUCTION

Channel systems are a family of distributed models where concurrent agents communicate via usually unbounded fifo communication buffers called “channels.” An agent of a channel system is modeled by a finite-state controller, and when taking a transition an agent can read messages from the channel or write into it. These models have turned out to be well-suited for the formal specification and algorithmic analysis of communication protocols and concurrent programs [Pac87, BG99, BH99, CF05, Mus10]. They are also a fundamental model of computation, closely related to Post’s tag systems. In all generality, channel systems are a Turing powerful model, which implies that most of their decision problems are undecidable.

A particularly interesting decidable and widely studied class of channel systems are the so-called *lossy channel systems* (LCSs), where channels are unreliable and may lose messages, see *e.g.* [CFP96, AJ96, BMO⁺12]. For LCSs, several important behavioral properties such as safety or inevitability are decidable. This is because, due to the lossy behavior of their channels, these systems are *well-structured*: transitions are monotonic with respect to a decidable well-quasi-ordering of the configuration space [AČJT00, FS01, SS13]. Beyond their applications in verification, LCSs have turned out to be an important automata-theoretic tool for decidability or hardness in areas like Timed Automata, Metric Temporal Logic,

2012 ACM CCS: [Theory of computation]: Models of computation; Semantics and reasoning.

Key words and phrases: Well quasi order; well-structured transition systems; fast-growing complexity.

* An extended abstract of this work first appeared in the proceedings of the 24th International Conference on Concurrency Theory (Concur’13) [HSS13].

Work partially funded by the REACHARD project ANR 11 BS02 001 01.

modal logics, *e.g.* [ADOW05, Kur06, OW07, LW08]. Moreover, they are also a fundamental model of computation capturing the \mathbf{F}_{ω} -complexity level in the fast-growing complexity hierarchy [Sch13], see [CS08, SS11].

Lossy channel systems do not provide an adequate way to model systems or protocols that treat messages discriminatingly according to some specified rule set. An example is the prioritization of messages, which is central to ensuring *quality of service* (QoS) properties in networking architectures, and is usually implemented by allowing for tagging messages with some relative priority. For instance, the Differentiated Services (DiffServ) architecture described in RFC 2475 [BBC⁺98], which enables QoS on modern IP networks, allows for a field specifying the relative priority of an IP packet with respect to a finite set of priorities, and network links may decide to arbitrarily drop IP packets of lower priority in favor of higher priority packets once the network congestion reaches a critical point. Another example of a similar priority-based policy arises in the context of ATM networks, where priorities are expressed via a single Cell Loss Priority bit in order to allow for giving preference (by dropping low-priority packages) to audio or video over less time-critical data [LB92].

Inspired by the aforementioned types of protocols, in this paper we introduce *priority channel systems* (PCSs), a family of channel systems where each message is equipped with a priority level, and where higher-priority messages can supersede lower-priority messages by dropping them. Priority channel systems rely on the *prioritized superseding ordering*, a novel ordering that generalizes Higman’s subword ordering and has not been considered before in the area of well-structured systems. It is however closely related to the gap-embedding considered in [SS85]. Showing it to be a well-quasi-ordering entails, among others, showing the decidability of safety and termination for PCSs. We complement our decidability results by showing that these problems become undecidable for channel systems that build upon more restrictive priority mechanisms, supporting the design choices made for our model.

1.1. Structure of this Paper. This paper can roughly be divided into two parts. In the first part, we define priority channel systems, explore this new model and analyze its power in complexity-theoretical terms. Beginning in Section 6, the second part relates priority channel systems in the broadest sense to related models or mathematical objects found in the literature.

In more detail, in Section 2 we provide an *at-a-glance* introduction to a simplified model of priority channel systems. This allows us to discuss on a high level the ideas behind our model, the main theorems, and the main algorithmic problems that we consider in this paper. We outline the decidability of fundamental decision problems via the framework of well-structured systems. Section 3 is then devoted to proving well-quasi-ordering properties of the prioritized superseding ordering which underlies priority channel systems. To this end, we characterize the superseding ordering via *priority embeddings*, which is an analogue and can in fact be seen as a generalization of Higman’s subword embedding. Using techniques from [SS11, SS85], we show in Section 4 an $\mathbf{F}_{\varepsilon_0}$ upper bound on the complexity of PCS verification, far higher than the \mathbf{F}_{ω} -complete complexity known for LCSs. We then prove in Section 5 a matching lower bound and this is the main technical result for PCSs of this paper: building upon techniques developed for less powerful models [CS08, Sch10b, HSS12], we show how PCSs can robustly simulate the computation of the fast growing functions F_α and their inverses for all ordinals α up to ε_0 . This gives a precise measure of the expressive power of PCSs.

In the second part of the paper, we first show in Section 6 that other natural choices of models of channel systems with priority mechanisms different from ours lead to undecidability of verification problems. We then turn to lossy channel systems and show in Section 7 that, although PCSs are not an extension of LCSs, they can very easily simulate LCS computations. In fact, we show how *higher-order* LCS models, which generalize the dynamic LCS from [AAC12], also embed into PCSs; this uses a more involved encoding of higher-order configurations and rules. Applications of the priority embedding to other well-quasi-ordered data structures such as depth-bounded trees found in the literature are subsequently discussed in Section 8.

2. PRIORITY CHANNEL SYSTEMS

In this section, we formally introduce Priority Channel Systems and give an overview about the decision problems we consider in this paper.

Definition 2.1. For every $d \in \mathbb{N}$, the *level- d priority alphabet* is $\Sigma_d \stackrel{\text{def}}{=} \{0, 1, \dots, d\}$. A *level- d priority channel system* (d -PCS) is a tuple $S = (\Sigma_d, \text{Ch}, Q, \Delta)$, where Σ_d is as above, $\text{Ch} = \{c_1, \dots, c_m\}$ is a set of m *channel names*, $Q = \{q_1, q_2, \dots\}$ is a finite set of *control states*, and $\Delta \subseteq Q \times \text{Ch} \times \{!, ?\} \times \Sigma_d \times Q$ is a set of *transition rules*.




FIGURE 1. A simple single-channel 3-PCS.

For the sake of a simplified introduction to PCSs in this section, the alphabet of a PCS abstracts away from actual message contents and only consists of natural numbers that indicate the priority of a message, where d is a message of highest and 0 of lowest priority, respectively. A treatment of more general alphabets is deferred to Section 3. The simple alphabet introduced here is however sufficient in order to show the lower bounds in Section 5. Moreover, from our definition it follows that Priority Channel Systems consist of a single process, which is sufficient for our purposes in this paper, since systems made of several concurrent components can be represented by a single process obtained as an asynchronous product of the components.

Figure 1 depicts a 3-PCS with a single channel and control states p and q . (A possible configuration of the channel is depicted alongside.) Informally speaking, when in control state q the PCS in Figure 1 can non-deterministically loop while writing the alphabet symbol “3” to the channel (to its right end), or switch to control state q if “3” can be read from the channel (from its left end). The key feature of PCSs is that messages with higher priority can erase messages with lower priority, *cf.* the formal semantics given next.

2.1. Semantics. The operational semantics of a PCS $S = (\Sigma_d, \text{Ch}, Q, \Delta)$ is given in terms of a transition system. We let $\text{Conf}_S \stackrel{\text{def}}{=} Q \times (\Sigma_d^*)^m$ be the set of all *configurations* of S , denoted C, D, \dots in the following. A configuration $C = (q, x_1, \dots, x_m)$ records an instantaneous control state $q \in Q$ and the contents of the m channels, *i.e.*, sequences of messages from Σ_d .

A sequence $x \in \Sigma_d^*$ has the form $x = a_1 \cdots a_\ell$ and we let $|x| = \ell$. Concatenation is denoted multiplicatively, with ε denoting the empty sequence.

The labeled transition relation between configurations, denoted $C \xrightarrow{\delta} C'$, is generated by the rules in $\Delta = \{\delta_1, \dots, \delta_k\}$. From a technical perspective, it is convenient to define two such transition relations, denoted \rightarrow_{rel} and $\rightarrow_{\#}$.

2.1.1. Reliable Semantics. We start with \rightarrow_{rel} that corresponds to “reliable” steps, or more correctly steps with no superseding of lower-priority messages. As is standard, for a *reading rule* of the form $\delta = (q, c_i?a, q') \in \Delta$, there is a step $C \xrightarrow{\delta}_{\text{rel}} C'$ if $C = (q, x_1, \dots, x_m)$ and $C' = (q', y_1, \dots, y_m)$ for some $x_1, y_1, \dots, x_m, y_m$ such that $x_i = a y_i$ and $x_j = y_j$ for all $j \neq i$, while for a *writing rule* $\delta = (q, c_i!a, q') \in \Delta$, there is a step $C \xrightarrow{\delta}_{\text{rel}} C'$ if $y_i = x_i a$ and $x_j = y_j$ for all $j \neq i$. These reliable steps correspond to the behavior of queue automata, or (reliable) channel systems, a Turing-powerful computation model [BZ83].

2.1.2. Internal-Superseding. The actual behavior of PCSs is obtained by extending reliable steps with *internal superseding steps*, denoted $C \xrightarrow{c_i\#k}_{\#} C'$, which can be performed at any time in an uncontrolled manner. Formally, for two words $x, y \in \Sigma_d^*$ and $k > 0$ in \mathbb{N} , we write

$$x \xrightarrow{\#k}_{\#} y \stackrel{\text{def}}{\iff}$$

- (1) x can be decomposed as $a_1 \cdots a_\ell$ with $\ell > k$;
- (2) $a_k \leq a_{k+1}$; and
- (3) $y = a_1 \cdots a_{k-1} a_{k+1} \cdots a_\ell$.

In other words, the k th message in x is superseded by its immediate successor a_{k+1} , with the condition that a_k is not of higher priority. We write $x \rightarrow_{\#} y$ when $x \xrightarrow{\#k}_{\#} y$ for some k , and use $x \leftarrow_{\#} y$ when $y \rightarrow_{\#} x$. The transitive reflexive closure $\leftarrow_{\#}^*$ is called the *superseding ordering* and is denoted by $\leq_{\#}$. Put differently, $\rightarrow_{\#}$ is a rewrite relation over Σ_d^* defined by the following string rewriting system (see [BJW82]):

$$\{a a' \rightarrow a' \mid 0 \leq a \leq a' \leq d\}. \quad (2.1)$$

This is extended to steps between configurations by $C \xrightarrow{c_i\#k}_{\#} C' \stackrel{\text{def}}{\iff} C = (q, x_1, \dots, x_m)$, $C' = (q', y_1, \dots, y_m)$, $q = q'$, $x_i \xrightarrow{\#k}_{\#} y_i$, and $x_j = y_j$ for $j \neq i$. Furthermore, every reliable step is a valid step: for any rule δ , $C \xrightarrow{\delta}_{\#} C'$ if $C \xrightarrow{\delta}_{\text{rel}} C'$, giving rise to a second transition system associated with S : $\mathcal{S}_{\#} \stackrel{\text{def}}{=} (\text{Conf}_S, \rightarrow_{\#})$.

Example 2.2. The following is a valid path in the transition system induced by the PCS from Figure 1:

$$(p, 0300) \xrightarrow{11}_{\#} (q, 03\underline{0}01) \xrightarrow{\#3}_{\#} (q, \underline{0}301) \xrightarrow{\#1}_{\#} (q, 3\underline{0}1) \xrightarrow{\#2}_{\#} (q, 31).$$

Here, underlining is used to show which symbol, if any, is superseded in the next step.

Remark 2.3 (Lossy Channel Systems). Priority channel systems and lossy channel systems are unreliable in the sense that uncontrolled rewrites may occur inside the channels. In the case of LCSs, arbitrary message losses can take place. The semantics of the two classes of systems thus differ quite a bit, however PCSs are easily seen to be able to encode LCS computations (by interspersing channel contents with a higher-priority symbol; see Section 7 and in particular Figure 10).

2.1.3. *Write-Superseding Semantics.* The internal-superseding semantics allows superseding to occur at any time and anywhere in the channel. Another possible scenario considers communications going through relays, network switches, or buffers, which handle incoming traffic with a so-called *write-superseding* policy, where writes immediately supersede (*i.e.*, erase) the congested messages in front of them. We develop this aspect here and prove the two semantics to be essentially equivalent.

Let $S = (\Sigma_d, \mathbf{Ch}, Q, \Delta)$ be a d -PCS. We define a new transition relation, denoted \rightarrow_w , between the configurations of S , giving rise to a transition system $\mathcal{S}_w \stackrel{\text{def}}{=} (\text{Conf}_S, \rightarrow_w)$. The relation \rightarrow_w is a variant of $\rightarrow_{\#}$ obtained by modifying the semantics of writing rules. Formally, for $\delta = (q, c_i!a, q') \in \Delta$, and for two configurations $C = (q, x_1, \dots, x_m)$ and $C' = (q, y_1, \dots, y_m)$, there is a step $C \xrightarrow{\delta}_w C'$ if $x_j = y_j$ for all $j \neq i$ and $y_i = za$ for a factorization $x_i = zz'$ of x_i where $z' \in \Sigma_a^*$, *i.e.*, where z' only contains messages from the level- a priority subalphabet. In other words, after $c_i!a$, the channel will contain a sequence y_i obtained from x_i by appending a in a way that may drop (erase) any number of suffix messages with priority $\leq a$, hence the “ $z' \in \Sigma_a^*$ ” requirement. The semantics of reading rules is unchanged so that $C \xrightarrow{\delta}_{\text{rel}} C'$ implies $C \xrightarrow{\delta}_w C'$.

Example 2.4. The PCS from Figure 1 has the following write-superseding run:

$$(p, 0 \underline{3} 0 0) \xrightarrow{!1}_w (q, 0 \underline{3} \underline{1}) \xrightarrow{!3}_w (q, 0 3) \xrightarrow{!3}_w (q, \underline{0} \underline{3} \underline{3}) \xrightarrow{!3}_w (q, 3) \xrightarrow{?3}_w (p, \varepsilon)$$

where in every configuration we underline the messages that will be superseded in the next step (and where, for simplicity, we do not write the full rule δ on the steps). Observe that $(p, 0 \underline{3} 0 0) \xrightarrow{*}_w (q, 3 1)$, to be contrasted with the internal-superseding run $(p, 0 \underline{3} 0 0) \xrightarrow{*}_{\#} (q, 3 1)$ in Example 2.2. Under write-superseding, the occurrence of 3 that is initially in the channel is not allowed to erase the 0 in front of it. \square

Compared to our standard PCS semantics, the write-superseding semantics adopts a localized viewpoint where the protocol managing priority levels and handling congestions resides at the sender’s end, and is not distributed all along the channels.

In the rest of this subsection, we show that the write-superseding is essentially equivalent to the standard semantics, see Proposition 2.5. A consequence is that one can freely choose to adopt either $\mathcal{S}_{\#}$ or \mathcal{S}_w as their favorite operational semantics for priority channel systems. In practice, we find it simpler to design and prove the correctness of some PCS—as we will in Sections 5 and 7—when assuming the write-superseding semantics since it is less liberal and easier to control. And we find it simpler to develop the formal theory of PCSs when assuming the internal-superseding semantics since it is finer-grained.

Proposition 2.5. *Let $C_0 = (q, \varepsilon, \dots, \varepsilon)$ be a configuration with empty channels, and C_f be any configuration. Then $C_0 \xrightarrow{\pm}_w C_f$ if, and only if, $C_0 \xrightarrow{\pm}_{\#} C_f$.*

The proof is organized in the three Lemmata 2.7–2.9 below.

Remark 2.6. Observe that the requirement of empty channels for C_0 in Proposition 2.5 cannot be lifted, as illustrated with Example 2.4. However, using standard coding tricks (*e.g.*, storing *initial* channel contents in control states), one can reduce a reachability or termination problem starting from an arbitrary initial configuration to the same question starting from an empty-channel C_0 , and show its decidability by combining Proposition 2.5 and Theorem 2.13. \square

Lemma 2.7 (From \mathcal{S}_w to $\mathcal{S}_\#$). *If \mathcal{S}_w has a run $C \xrightarrow{+}_w D$ then $\mathcal{S}_\#$ has a run $C \xrightarrow{+}_\# D$.*

Proof. We show that $\xrightarrow{+}_w$ is contained in $\xrightarrow{+}_\#$, assuming for the sake of simplicity that S has only one channel.

A writing step $(p, x) \xrightarrow{!a}_w (q, y)$ with $x = z b_1 \cdots b_j$ and $y = z a$ in \mathcal{S}_w can be simulated in $\mathcal{S}_\#$ with $(p, x) \xrightarrow{!a}_\# (q, z b_1 \cdots b_j a) \xrightarrow{\# \ell}_\# (q, z b_1 b_{j-1}) \xrightarrow{\# \ell-1}_\# \cdots \xrightarrow{\# k+1}_\# (q, z a)$, where $\ell = |x|$ and $k = |z|$. Reading steps simply coincide in \mathcal{S}_w and $\mathcal{S}_\#$. \square

In the other direction, one can translate runs in $\mathcal{S}_\#$ to runs in \mathcal{S}_w as stated by following lemma.

Lemma 2.8 (From $\mathcal{S}_\#$ to \mathcal{S}_w). *If $\mathcal{S}_\#$ has a run $C \xrightarrow{*}_\# D$ then \mathcal{S}_w has a run $C' \xrightarrow{*}_w D$ for some $C' \leq_\# C$. In particular, if the channels are empty in C , then necessarily $C' = C$ and $C \xrightarrow{*}_w D$.*

Proof. Again we assume that S has only one channel.

Write the run $C \xrightarrow{*}_\# D$ under the form $C_0 \rightarrow_\# C_1 \rightarrow_\# \cdots \rightarrow_\# C_n$ and rearrange its steps so that superseding occurs greedily. This relies on Lemma 2.9 stated next without proof.

Repeatedly applying Lemma 2.9 to transform $C_0 \xrightarrow{*}_\# C_n$ as long as possible is bound to terminate (with each commutation, superseding steps are shifted to the left of reliable steps, or the sum $\sum_i k_i$ of superseding positions in steps $C_{i-1} \xrightarrow{\# k_i}_\# C_i$ increases strictly while being bounded by $O(n^2)$ for a length- n run). One eventually obtains a new run $C_0 \xrightarrow{*}_\# C_n$ with same starting and final configurations, and where all the superseding steps occur (at the beginning of the run or) just after a write in *normalized* sequences of the form

$$C = (q, x) \xrightarrow{!a}_\# \xrightarrow{\# \ell}_\# \xrightarrow{\# \ell-1}_\# \xrightarrow{\# \ell-2}_\# \cdots \xrightarrow{\# \ell-r}_\# C', \quad (2.2)$$

where furthermore $\ell = |x|$. In this case, \mathcal{S}_w has a step $C \xrightarrow{!a}_w C'$.

Greedily shifting superseding steps to the left may move some of them at the start of the run instead of after a write: these steps are translated into $C \geq_\# C'$ in Lemma 2.8. Finally, the steps that are not in normalized sequences are reading steps which exist unchanged in \mathcal{S}_w . \square

Lemma 2.9 (Commuting $\#$ -steps).

- (1) *If $C_1 \xrightarrow{?a}_\# C_2 \xrightarrow{\# k}_\# C_3$ then there is a configuration C'_2 s.t. $C_1 \xrightarrow{\# k+1}_\# C'_2 \xrightarrow{?a}_\# C''$.*
- (2) *If $C_1 = (q, x) \xrightarrow{!a}_\# C_2 \xrightarrow{\# k}_\# C_3$ with $k < |x|$, then there is a configuration C'_2 s.t. $C_1 \xrightarrow{\# k}_\# C'_2 \xrightarrow{!a}_\# C_3$.*
- (3) *If $C_1 = (q, x) \xrightarrow{\# k_1}_\# C_2 \xrightarrow{\# k_2}_\# C_3$ with $k_1 \leq k_2$ then there is a configuration C'_2 s.t. $C_1 \xrightarrow{\# k_2+1}_\# C'_2 \xrightarrow{\# k_1}_\# C''$.*

2.2. Priority Channel Systems are Well-Structured. Our main result regarding the verification of PCSs is that they are *well-structured* systems, which entails the decidability of standard decision problems via the generic decidability results from [AČJT00, FS01, SS13]. Let us first recall the definitions of well-quasi-orders and well-structured systems.

Definition 2.10 (wqo). Let (A, \leq_A) be a quasi order. Then (A, \leq_A) is a *well-quasi-order* (wqo) if for any infinite sequence x_0, x_1, x_2, \dots over A there exists two indices $i < j$ such that $x_i \leq_A x_j$.

A simple example of a wqo is any finite set Σ with equality $(\Sigma, =)$, thanks to the pigeonhole principle. More generally, complex wqos can be build from simpler ones by algebraic operations [SS12]. Let (A_1, \leq_{A_1}) and (A_2, \leq_{A_2}) be wqos:

- Their *disjoint sum* $A_1 + A_2 \stackrel{\text{def}}{=} \{\langle x, i \rangle \mid i \in \{1, 2\} \text{ and } x \in A_i\}$ is well-quasi-ordered by the *sum ordering* \leq_+ defined by $\langle x, i \rangle \leq_+ \langle y, j \rangle \stackrel{\text{def}}{\iff} i = j \text{ and } x \leq_{A_i} y$.
- Their *Cartesian product* $A_1 \times A_2 \stackrel{\text{def}}{=} \{\langle x, y \rangle \mid x \in A_1 \text{ and } y \in A_2\}$ is well-quasi-ordered by the *product ordering* \leq_\times defined by $\langle x, y \rangle \leq_\times \langle x', y' \rangle \stackrel{\text{def}}{\iff} x \leq_{A_1} x' \text{ and } y \leq_{A_2} y'$. This is also known as Dickson’s Lemma.
- The Kleene star A_1^* , *i.e.*, the set of finite sequences over A_1 is well-quasi-ordered by the *substring embedding* relation \leq_* defined by $x \leq_* y \stackrel{\text{def}}{\iff} x = a_1 \cdots a_\ell, y = y_0 b_1 y_1 \cdots y_\ell b_\ell y_{\ell+1}$ for some a_i, b_i in A , some y_i in A^* , and such that $a_i \leq_{A_1} b_i$ for every $1 \leq i \leq \ell$. This is known as Higman’s Lemma, and is instrumental in the study of lossy channel systems (*cf.* Section 6).

Definition 2.11 (WSTS). A *well-structured (transition) system* (WSTS) is a tuple $\mathcal{S} = (A, \rightarrow, \leq_A)$ with $\rightarrow \subseteq A \times A$ such that

- (1) (A, \leq_A) is a wqo; and
- (2) \rightarrow is compatible with respect to \leq_A , *i.e.*, if $x \rightarrow y$ and $x \leq_A x'$ then there is some y' such that $x' \xrightarrow{*} y'$ and $y \leq_A y'$.

A WSTS enjoys a stronger *stuttering compatibility* if the second condition is altered to require $x' \xrightarrow{\pm} y'$, see [FS01, Definition 4.4]. Let $S = (\Sigma_d, \text{Ch}, Q, \Delta)$ be a PCS, we define the following order on configurations of S : $C \leq_\# D \stackrel{\text{def}}{\iff} C$ is some (p, y_1, \dots, y_m) and D is (p, x_1, \dots, x_m) with $x_i \leq_\# y_i$ for all $i = 1, \dots, m$. Equivalently, $C \leq_\# D$ if C can be obtained from D by internal superseding steps.

Theorem 2.12 (PCSs are WSTSs). *For any PCS S , $\mathcal{S}_\# = (\text{Conf}_S, \rightarrow_\#, \leq_\#)$, *i.e.*, the transition system $\mathcal{S}_\#$ with configurations ordered by $\leq_\#$, is a well-structured system with stuttering compatibility.*

Proof. We have to show that the two conditions required in Definition 2.11 hold. Proving that $(\text{Conf}_S, \leq_\#)$ is a well-quasi-ordering is the topic of Section 3 and will be established in a more general setting in Theorem 3.6.

Checking stuttering compatibility is trivial with the $\leq_\#$ ordering. Indeed, assume that $C \leq_\# C'$ and that $C \rightarrow_\# D$ is a step from the “smaller” configuration. Then in particular $C' \xrightarrow{*}_\# C$ by definition of $\rightarrow_\#$, so that clearly $C' \xrightarrow{\pm}_\# D$ and C' can simulate any step from C . \square

A consequence of the well-structuredness of PCSs is the decidability of several natural verification problems. In this paper we focus on “Reachability,” aka “Safety” when we want to check that a configuration is *not* reachable: given a PCS, an initial configuration C_0 , and a recursive set of configurations $G \subseteq \text{Conf}_S$, does $C_0 \xrightarrow{*}_\# D$ for some $D \in G$? Another decision problem is “Inevitability,” *i.e.* to decide whether all maximal runs from C_0 eventually visit G , which includes “Termination” as a special case.

Theorem 2.13. *Reachability and Inevitability are decidable for PCSs.*

Proof (Sketch). In order to apply the generic WSTS algorithms from [FS01], we have to prove that the order $\leq_{\#}$ is decidable, and that the set of immediate successors of a configuration and the minimal immediate predecessors of an upward-closed set are computable.

Deciding the ordering $\leq_{\#}$ between configurations is in NLOGSPACE; the proof of this fact is the purpose of Remark 3.7. Moreover, the operational semantics is finitely branching and effective, *i.e.*, one can compute the immediate successors of a configuration and the minimal immediate predecessors of an upward-closed set.

We note that Reachability and Coverability coincide (even for zero-length runs when C_0 has empty channels) since $\overset{\pm}{\rightarrow}_{\#}$ coincides with $\geq_{\#} \circ \overset{\pm}{\rightarrow}_{\#}$, and that the answer to a Reachability question only depends on the (finitely many) minimal elements of G . One can even compute $Pre^*(G)$ for G given, *e.g.*, as a regular subset of $Conf_G$.

For Inevitability, the algorithms in [AČJT00, FS01] assume that G is downward-closed but, in our case where $\overset{\pm}{\rightarrow}_{\#}$ and $\geq_{\#} \circ \overset{\pm}{\rightarrow}_{\#}$ coincide, decidability can be shown for arbitrary (recursive) G , as in [Sch10a, Theorem 4.4]. \square

3. PRIORITY EMBEDDING

In this section we establish that the superseding ordering $\leq_{\#}$ on words enjoys the well-quasi-ordering properties we require for reasoning about PCSs. In order to keep our results generic, as already stated at the beginning of Section 2, we establish those properties over an alphabet that is more general than the one introduced in Definition 2.1. Instead of allowing for messages consisting merely of priorities, we allow for messages over an arbitrary well-quasi-ordering to be tagged with priority numbers. This is in line with the algebraic operations on wqos presented at the beginning of Section 2.2.

Definition 3.1 (Generalized Priority Alphabet). Let $d \in \mathbb{N}$ be a *priority level* and let (Γ, \leq_{Γ}) be a well-quasi-order, a *generalized level- d priority alphabet over Γ* is $\Sigma_{d,\Gamma} \stackrel{\text{def}}{=} \{(a, w) \mid 0 \leq a \leq d, w \in \Gamma\}$.

Subsequently, we call $\Sigma_{d,\Gamma}$ a *generalized priority alphabet* for brevity. In analogy to the internal superseding steps in Section 2.1, we define the *generalized priority relation* $\rightarrow_{\#, \Gamma}$ over finite strings in $\Sigma_{d,\Gamma}^*$ via a string rewriting system with the following two families of rule schemata:

$$\{(a, w)(a', w') \rightarrow_{\#, \Gamma} (a', w') \mid a \leq a', w \in \Gamma\}, \quad (3.1)$$

$$\{(a, w) \rightarrow_{\#, \Gamma} (a, w') \mid w' \leq_{\Gamma} w\}. \quad (3.2)$$

Informally speaking, the first line states that a string can be rewritten if some higher-priority message supersedes a lower priority message, and the second that any message can be rewritten to a message that is below in the wqo (Γ, \leq_{Γ}) .

We define $\leq_{\#, \Gamma} \stackrel{\text{def}}{=} \overset{*}{\leftarrow}_{\#}$, *i.e.*, $\leq_{\#, \Gamma}$ is the reflexive transitive closure of the inverse of $\rightarrow_{\#, \Gamma}$. The main purpose of this section is to prove that $(\Sigma_{d,\Gamma}^*, \leq_{\#, \Gamma})$ is a well-quasi-ordering, *cf.* Definition 2.11. To this end, we will first establish a characterization of $\leq_{\#, \Gamma}$ via an embedding relation and subsequently prove that the obtained priority embeddings yield a well-quasi-ordering.

Before we continue, let us remark that the priority alphabet in Definition 2.1 and the relation $\rightarrow_{\#}$ from Section 2.1 are obtained by considering $(\Gamma, =)$ for some singleton set Γ . Whenever we drop the index Γ , we implicitly refer to this well-quasi-ordering. Letting Γ be a finite set of messages represented as strings and \leq_{Γ} the identity relation yields a generalized priority alphabet where a priority can be assigned to each message. Such an alphabet underlies for instance the well-quasi-ordering that we will later use for showing that planar planted trees are well-quasi-ordered under minors, *cf.* Section 8.2. Another example is $\Gamma = \Sigma^*$ for some finite alphabet Σ and where \leq_{Γ} is the substring embedding, which allows for representing unbounded messages on a lossy channel which are tagged with a priority level. Finally, we wish to mention that for a generalized priority alphabet $\Sigma_{d,\Gamma}$, if we wish to apply $\Sigma_{d,\Gamma}$ in a PCSs, for Theorem 2.13 to hold, (Γ, \leq_{Γ}) has to fulfill the same properties: \leq_{Γ} should be decidable, and the set of immediate successors of a configuration and the minimal immediate predecessors of an upward-closed set should be computable.

3.1. Embedding with Priorities. We define here an embedding relation between finite strings over a generalized priority alphabet $\Sigma_{d,\Gamma}^*$. We are inspired in this by a coarser *gap embedding* relation defined by Schütte and Simpson [SS85], who use it to derive a “natural” formal statement undecidable in Peano arithmetic—the totality of the function H^{ε_0} defined later in Section 4 is another well-known example of such a statement. Besides refining the gap embedding relation to match the superseding ordering, we also extend it to handle an underlying wqo (Γ, \leq_{Γ}) .

Given $x, y \in \Sigma_{d,\Gamma}^*$, we define the *generalized priority embedding* $\sqsubseteq_{p,\Gamma}$ by

$$x \sqsubseteq_{p,\Gamma} y \stackrel{\text{def}}{\iff} \begin{cases} x \text{ is some } (a_1, v_1) \cdots (a_{\ell}, v_{\ell}) \\ y \text{ is some } y_1 (a_1, w_1) y_2 (a_2, w_2) \cdots y_{\ell} (a_{\ell}, w_{\ell}) \\ \text{such that } \forall i = 1, \dots, \ell : y_i \in \Sigma_{a_i,\Gamma}^* \text{ and } v_i \leq_{\Gamma} w_i . \end{cases}$$

For example, in the singleton case, $201 \sqsubseteq_p 22011$ but $120 \not\sqsubseteq_p 10210$, since factoring 10210 as $z_1 1 z_2 2 z_3 0$ would require $z_3 = 1 \notin \Sigma_0^*$. If $x \sqsubseteq_p y$ then x is a subword of y and x can be obtained from y by removing factors of messages with priority not above the first preserved message to the right of the factor. Observe that $\sqsubseteq_{p,\Gamma}$ is similar (but not equivalent) to the Higman subword embedding for $d = 0$. From the definition above, we get the following properties which we will implicitly use subsequently:

$$\varepsilon \sqsubseteq_{p,\Gamma} y \quad \text{iff} \quad y = \varepsilon , \quad (3.3)$$

$$x_1 \sqsubseteq_{p,\Gamma} y_1, x_2 \sqsubseteq_{p,\Gamma} y_2 \quad \text{imply} \quad x_1 x_2 \sqsubseteq_{p,\Gamma} y_1 y_2 , \quad (3.4)$$

$$x_1 x_2 \sqsubseteq_{p,\Gamma} y \quad \text{implies} \quad \exists y_1 \exists_{p,\Gamma} x_1 : \exists y_2 \exists_{p,\Gamma} x_2 : y = y_1 y_2 , \quad (3.5)$$

$$v \leq_{\Gamma} w \quad \text{implies} \quad \forall 0 \leq a \leq d : \forall z \in \Sigma_{a,\Gamma}^* : (a, v) \sqsubseteq_{p,\Gamma} z(a, w) . \quad (3.6)$$

Lemma 3.2. *Let $\Sigma_{d,\Gamma}$ be a generalized priority alphabet. Then $(\Sigma_{d,\Gamma}^*, \sqsubseteq_{p,\Gamma})$ is a quasi-ordering.*

Proof. We have to show that $(\Sigma_{d,\Gamma}^*, \sqsubseteq_{p,\Gamma})$ is reflexive and transitive. Reflexivity is obvious from the definition of $\sqsubseteq_{p,\Gamma}$. Regarding transitivity, let $x, y, z \in \Sigma_{d,\Gamma}^*$ be such that $x \sqsubseteq_{p,\Gamma} y \sqsubseteq_{p,\Gamma} z$ and write $x = (a_1, u_1) \cdots (a_{\ell}, u_{\ell})$. Since $x \sqsubseteq_{p,\Gamma} y$, by definition we can write $y = y_1 (a_1, v_1) \cdots y_{\ell} (a_{\ell}, v_{\ell})$, where $u_i \leq_{\Gamma} v_i$ and each $y_i = (b_{1,i}, v_{1,i}) \cdots (b_{m_i,i}, v_{m_i,i}) \in \Sigma_{a_i,\Gamma}^*$ for all $1 \leq i \leq \ell$. Consequently, since $y \sqsubseteq_{p,\Gamma} z$, we can decompose z as $z = z_1 (a_1, w_1) \cdots z_{\ell} (a_{\ell}, w_{\ell})$,

where each z_i is of the form

$$z_i = z_{1,i}(b_{1,i}, w_{1,i}) \cdots z_{m_i,i}(b_{m_i,i}, w_{m_i,i}) z'_i.$$

Since each $(b_{j,i}, w_{j,i}) \in \Sigma_{a_i, \Gamma}^*$, by definition of $\sqsubseteq_{p, \Gamma}$ we have $z_i \in \Sigma_{a_i, \Gamma}^*$, hence the above decomposition of z in particular yields $x \sqsubseteq_{p, \Gamma} z$. \square

The generalized priority embedding acts as a relational counterpart to the more operational generalized superseding ordering. In fact $\leq_{\#, \Gamma}$ and $\sqsubseteq_{p, \Gamma}$ coincide, as shown by the next lemma.

Lemma 3.3. *For any $x, y \in \Sigma_{d, \Gamma}^*$, $x \leq_{\#, \Gamma} y$ if, and only if, $x \sqsubseteq_{p, \Gamma} y$.*

Proof. In the following, write x as $x = (a_1, v_1) \cdots (a_k, v_k)$.

Suppose $x \leq_{\#, \Gamma} y$, i.e., $y \xrightarrow{*}_{\#, \Gamma} x$. We show $x \sqsubseteq_{p, \Gamma} y$ by induction on the number of superseding steps. The base case where no superseding occurs entails $x = y$ and we rely on the reflexivity of $\sqsubseteq_{p, \Gamma}$. For the induction step, let $y \rightarrow_{\#, \Gamma} z$ such that $z \xrightarrow{*}_{\#, \Gamma} x$. By the induction hypothesis, $x \sqsubseteq_{p, \Gamma} z$, i.e., z can be factored as $z = z_1(a_1, w_1) \cdots z_k(a_k, w_k)$ such that $z_i \in \Sigma_{a_i, \Gamma}^*$ and $v_i \leq_{\Gamma} w_i$ for all $1 \leq i \leq k$. We do a case distinction on which rewriting rule is applied in order to obtain $y \rightarrow_{\#, \Gamma} z$:

- If $y \rightarrow_{\#, \Gamma} z$ via (3.1) then y is obtained from z by replacing some $z_j = z_{j,1} \cdots z_{j,\ell_j}$ with $z'_j = z_{j,1} \cdots z_{j,i-1}(b, w) z_{j,i} \cdots z_{j,\ell_j}$ for some $1 \leq i \leq \ell_j + 1$, $1 \leq j \leq k$ and (b, w) such that in particular $b \leq a_j$, and hence $z'_j \in \Sigma_{a_j, \Gamma}^*$. Thus y factors as $y = z_1(a_1, w_1) \cdots z'_j(a_j, w_j) \cdots z_k(a_k, w_k)$, which allows us to conclude that $x \sqsubseteq_{p, \Gamma} y$.
- If $y \rightarrow_{\#, \Gamma} z$ via (3.2), y is obtained by replacing some (a, w) occurring in z with (a, w') for some $w' \geq_{\Gamma} w$. By transitivity of \leq_{Γ} , $x \sqsubseteq_{p, \Gamma} y$ follows immediately.

Conversely, assume $x \sqsubseteq_{p, \Gamma} y$: then y factors as $y = y_1(a_1, w_1) \cdots y_k(a_k, w_k)$. Since for every (a, w) occurring in some y_i we have $a \leq a_i$, by repeatedly applying (3.1) we have $y \xrightarrow{*}_{\#, \Gamma} z = (a_1, w_1) \cdots (a_k, w_k)$. Moreover, $v_i \leq_{\Gamma} w_i$ for all $1 \leq i \leq k$, and thus by repeated application of (3.2) we get $z \xrightarrow{*}_{\#, \Gamma} x$, as required. \square

3.2. Priority Embedding is a Well-Quasi-Ordering. The purpose of this section is to prove that $\sqsubseteq_{p, \Gamma}$ is a well-quasi-ordering. By application of Lemma 3.3, this entails that $\leq_{\#, \Gamma}$ is a well-quasi-ordering as well. We rely for this on the algebraic operations presented in Section 2.2, and on a classical tool from wqo theory, namely *order reflections*:

Definition 3.4 (Order Reflection). Let (A, \leq_A) and (B, \leq_B) be two quasi-orders. An *order reflection* is a mapping $r: A \rightarrow B$ such that $r(x) \leq_B r(y)$ implies $x \leq_A y$.

The following is folklore (and easy to see):

Fact 3.5. *Let (A, \leq_A) and (B, \leq_B) be two quasi-orders and r be an order reflection $A \rightarrow B$. If (B, \leq_B) is a wqo, then (A, \leq_A) is a wqo.*

In the following, we define the *height* of a sequence $x \in \Sigma_{d, \Gamma}^*$, written $h(x)$, as being the highest priority occurring in x . By convention, we let $h(\varepsilon) \stackrel{\text{def}}{=} -1$. Thus, $x \in \Sigma_{h, \Gamma}^*$ if and only if $h \geq h(x)$, and we further let $\Sigma_{-1, \Gamma} \stackrel{\text{def}}{=} \emptyset$. Any $x \in \Sigma_{d, \Gamma}^*$ has a unique *canonical* factorization $x = x_0(h, v_1)x_1 \cdots x_{m-1}(h, v_m)x_m$ where m is the number of occurrences of $h = h(x)$ in x and where the $m + 1$ *residuals* x_0, x_1, \dots, x_m are in $\Sigma_{h-1, \Gamma}^*$.

Theorem 3.6. *Let $\Sigma_{d,\Gamma}$ be a generalized priority alphabet. Then $(\Sigma_{d,\Gamma}^*, \sqsubseteq_{p,\Gamma})$ is a well-quasi-ordering.*

Proof. We proceed by induction on d . For the base case $d = -1$, i.e. for the empty priority alphabet, $(\Sigma_{-1,\Gamma}^*, \sqsubseteq_{p,\Gamma}) = (\{\varepsilon\}, =)$ is a wqo.

For the induction step, a word $x \in \Sigma_{d,\Gamma}^*$ is either in $\Sigma_{d-1,\Gamma}^*$, or it has height $h(x) = d$ and then the canonical height factoring lets us write x under the form

$$x = x_0(d, v_1)x_1 \cdots x_{m-1}(d, v_m)x_m \quad (3.7)$$

with residuals $x_i \in \Sigma_{d-1,\Gamma}^*$ for all $0 \leq i \leq m$. By the induction hypothesis, $(\Sigma_{d-1,\Gamma}^*, \sqsubseteq_{p,\Gamma})$ is a well-quasi-ordering. We exhibit an order reflection $r: \Sigma_{d,\Gamma}^* \rightarrow \Theta_{d,\Gamma}$, where

$$\Theta_{d,\Gamma} \stackrel{\text{def}}{=} \Sigma_{d-1,\Gamma}^* + \Sigma_{d-1,\Gamma}^* \times ((\{d\} \times \Gamma) \times \Sigma_{d-1,\Gamma}^*)^* \times (\{d\} \times \Gamma) \times \Sigma_{d-1,\Gamma}^*. \quad (3.8)$$

Since $\Theta_{d,\Gamma}$ is obtained from the well-quasi-orders $(\Sigma_{d-1,\Gamma}^*, \sqsubseteq_{p,\Gamma})$, (Γ, \leq_Γ) , and $(\{d\}, =)$ by disjoint sum, Cartesian product, and substring embedding, Fact 3.5 will allow us to conclude that $(\Sigma_{d,\Gamma}^*, \sqsubseteq_{p,\Gamma})$ is a well-quasi-order. To this end, for x and m as above, if $m > 0$ we define

$$r(x) \stackrel{\text{def}}{=} (x_0, [((d, v_1), x_1) \cdots ((d, v_{m-1}), x_{m-1})], (d, v_m), x_m), \quad (3.9)$$

and $r(x) \stackrel{\text{def}}{=} x_0 = x$ if $m = 0$. We need to verify that, whenever $r(x) \preceq r(y)$ with respect to the ordering \preceq associated with $\Theta_{d,\Gamma}$ by the algebraic operations, then $x \sqsubseteq_{p,\Gamma} y$. This is obvious when both $r(x) = x$ and $r(y) = y$ are in $\Sigma_{d-1,\Gamma}^*$. Otherwise, let $r(x)$ be as in (3.9) and write

$$r(y) = (y_0, [((d, w_1), y_1) \cdots ((d, w_{n-1}), y_{n-1})], (d, w_n), y_n).$$

Since $r(x) \preceq r(y)$, from the product ordering we obtain

$$x_0 \sqsubseteq_{p,\Gamma} y_0, \quad (3.10)$$

$v_m \leq_\Gamma w_n$, and $x_m \sqsubseteq_{p,\Gamma} y_n$, while from the subword ordering we obtain the existence of indices $1 \leq i_1, \dots, i_{m-1} < n$ such that $v_j \leq_\Gamma w_{i_j}$ and $x_j \sqsubseteq_{p,\Gamma} y_{i_j}$ for all $0 < j < m$. Setting $i_0 \stackrel{\text{def}}{=} 0$ and $i_m \stackrel{\text{def}}{=} n$, observe that by (3.4) and (3.6),

$$(d, v_j)x_j \sqsubseteq_{p,\Gamma} (d, w_{i_{j-1}+1})y_{i_{j-1}+1} \cdots y_{i_j-1}(d, w_{i_j})y_{i_j} \quad (3.11)$$

for all $0 < j \leq m$, which together with (3.10) and (3.4) implies $x \sqsubseteq_{p,\Gamma} y$ as desired. \square

Remark 3.7. Theorem 3.6 and Lemma 3.3 prove that $\leq_\#$ is a wqo on configurations of PCSs, as we assumed in Section 2.2. There we also assumed that $\leq_\#$ is decidable. We can now see that it is in NLOGSPACE, since, in view of Lemma 3.3, one can check whether $x \leq_\# y$ by reading x and y simultaneously while guessing nondeterministically a factorization $z_1 a_1 \cdots z_\ell a_\ell$ of y , and checking that $z_i \in \Sigma_{a_i}^*$. \square

4. FAST-GROWING UPPER BOUNDS

The verification of infinite-state systems, and WSTSs in particular, often turns out to require astronomic computational resources expressed as *subrecursive functions* [LW70, FW98] of the input size. We show in this section how to bound the complexity of the algorithms presented in Section 2.2 and classify the Reachability and Inevitability problems for PCSs using *fast-growing complexity classes* [Sch13].

To this end, we first provide the necessary background on subrecursive functions in Section 4.1. The heart of the upper bound proof is a specialized Length Function Theorem for $(\Sigma_{d,\Gamma}^*, \sqsubseteq_{p,\Gamma})$, obtained in Section 4.2 by instrumenting the proof of Theorem 3.6 and applying the generic Length Function Theorem from [SS11]. This allows us to derive $\mathbf{F}_{\varepsilon_0}$ upper bounds and new *combinatorial algorithms* for PCS verification in Section 4.3.

4.1. Subrecursive Hierarchies. Throughout this paper, we use *ordinal terms* inductively defined by the following grammar

$$(\Omega \ni) \alpha, \beta, \gamma ::= 0 \mid \omega^\alpha \mid \alpha + \beta$$

where addition is associative, with 0 as the neutral element (the empty sum). Such a term $\alpha = \sum_{i=0}^k \omega^{\alpha_i}$ is 0 if $k = 0$, otherwise a *successor* if $\alpha_k = 0$ and a *limit* otherwise. We often write 1 as short-hand for ω^0 , and ω for ω^1 . The symbol λ is reserved for limit ordinal terms.

We can associate a set-theoretic ordinal $o(\alpha)$ with each term α by interpreting $+$ as the direct sum operator and ω as \mathbb{N} ; this gives rise to a well-founded quasi-ordering $\alpha < \beta \stackrel{\text{def}}{\iff} o(\alpha) < o(\beta)$. A term $\alpha = \sum_{i=1}^k \omega^{\alpha_i}$ is in *Cantor normal form* (CNF) if $\alpha_1 \geq \alpha_2 \geq \dots \geq \alpha_k$ and each α_i is itself in CNF for $i = 1, \dots, k$. Terms in CNF and set-theoretic ordinals below ε_0 are in bijection; it will however be convenient later in Section 5 to consider terms that are *not* in CNF.

With any limit term λ , we associate a *fundamental sequence* of terms $(\lambda_n)_{n \in \mathbb{N}}$:

$$(\gamma + \omega^{\beta+1})_n \stackrel{\text{def}}{=} \gamma + \omega^\beta \cdot n = \gamma + \overbrace{\omega^\beta + \dots + \omega^\beta}^n, \quad (\gamma + \omega^\lambda)_n \stackrel{\text{def}}{=} \gamma + \omega^{\lambda'_n}. \quad (4.1)$$

This yields $\lambda_0 < \lambda_1 < \dots < \lambda$ for any λ , with furthermore $\lambda = \lim_{n \in \mathbb{N}} \lambda_n$. For instance, $\omega_n = n$, $(\omega^\omega)_n = \omega^n$, etc. Note that λ_n is in CNF when λ is.

We need to add a term ε_0 to Ω to represent the set-theoretic ε_0 , *i.e.*, the smallest solution of $x = \omega^x$. We take this term to be a limit term as well; we define the fundamental sequence for ε_0 by $(\varepsilon_0)_n \stackrel{\text{def}}{=} \Omega_n$, where for $n \in \mathbb{N}$, we use Ω_n as short-hand notation for the ordinal $\omega^{\omega^{\dots^\omega}} \}$ n stacked ω 's, *i.e.*, for $\Omega_0 \stackrel{\text{def}}{=} 1$ and $\Omega_{n+1} \stackrel{\text{def}}{=} \omega^{\Omega_n}$.

4.1.1. Inner Recursion Hierarchies. Our main subrecursive hierarchy is the *Hardy hierarchy*. Given a monotone expansive unary function $h: \mathbb{N} \rightarrow \mathbb{N}$, it is defined as an ordinal-indexed hierarchy of unary functions $(h^\alpha: \mathbb{N} \rightarrow \mathbb{N})_\alpha$ through

$$h^0(n) \stackrel{\text{def}}{=} n, \quad h^{\alpha+1}(n) \stackrel{\text{def}}{=} h^\alpha(h(n)), \quad h^\lambda(n) \stackrel{\text{def}}{=} h^{\lambda_n}(n). \quad (4.2)$$

Observe that h^1 is simply h , and more generally h^α is the α th iterate of h , using diagonalization to treat limit ordinals.

A case of particular interest is to choose the successor function $H(n) \stackrel{\text{def}}{=} n + 1$ for h . Then the *fast growing hierarchy* $(F_\alpha)_\alpha$ can be defined by $F_\alpha \stackrel{\text{def}}{=} H^{\omega^\alpha}$, resulting in

$F_0(n) = H^1(n) = n + 1$, $F_1(n) = H^\omega(n) = H^n(n) = 2n$, $F_2(n) = H^{\omega^2}(n) = 2^n n$ being exponential, $F_3 = H^{\omega^3}$ being non-elementary, $F_\omega = H^{\omega^\omega}$ being an Ackermannian function, F_{ω^k} a k -Ackermannian function, and $F_{\varepsilon_0} = H^{\varepsilon_0} \circ H$ a function whose totality is not provable in Peano arithmetic [FW98].

4.1.2. Fast-Growing Complexity Classes. Our intention is to establish the “ F_{ε_0} completeness” of verification problems on PCSs. In order to make this statement more precise, we define the class $\mathbf{F}_{\varepsilon_0}$ as a specific instance of the *fast-growing complexity classes* defined for $\alpha \geq 3$ by [Sch13]

$$\mathbf{F}_\alpha \stackrel{\text{def}}{=} \bigcup_{p \in \bigcup_{\beta < \alpha} \mathcal{F}_\beta} \text{DTIME}(F_\alpha(p(n))), \quad \mathcal{F}_\alpha = \bigcup_{c < \omega} \text{FDTIME}(F_\alpha^c(n)), \quad (4.3)$$

where the class of *functions* \mathcal{F}_α as defined above is the α th level of the *extended Grzegorzcyk hierarchy* [LW70] when $\alpha \geq 2$.

The latter hierarchy of function classes $(\mathcal{F}_\alpha)_\alpha$ is well-established [LW70, FW98]. The class \mathcal{F}_α is the set of functions computable in time F_α^c , a finite iterate of F_α . In particular, \mathcal{F}_2 is the set of elementary functions, $\bigcup_{\alpha < \omega} \mathcal{F}_\alpha$ the set of primitive-recursive functions, while $\bigcup_{\alpha < \varepsilon_0} \mathcal{F}_\alpha$ is exactly the set of ordinal-recursive (aka “provably recursive”) functions [FW98].

The complexity classes $(\mathbf{F}_\alpha)_\alpha$ are more recent [Sch13]: \mathbf{F}_α is the set of decision problems that can be solved in time $F_\alpha \circ p$ for some p in $\bigcup_{\beta < \alpha} \mathcal{F}_\beta$. Each \mathbf{F}_α is naturally equipped with $\bigcup_{\beta < \alpha} \mathcal{F}_\beta$ as classes of reductions. For instance, \mathcal{F}_2 is the set of elementary functions, and \mathbf{F}_3 the class of problems whose complexity is bounded by a tower of exponentials of height given by some elementary function of the input.¹

4.2. The Length of Controlled Bad Sequences. A finite or infinite sequence x_0, x_1, \dots over a quasi-order (A, \leq_A) is called *bad* if, for all indices $i < j$, $x_i \not\leq_A x_j$. Definition 2.10 can thus be restated by saying that (A, \leq_A) is a wqo if and only if every bad sequence over A is finite. In order to bound the complexity of the algorithms from Theorem 2.13, we wish to bound the lengths of bad sequences over the wqo $(\text{Conf}_S, \leq_\#)$. More precisely, the main issue here is to bound the length of bad sequences over $(\Sigma_d^*, \sqsubseteq_p)$; we actually work in the more general case of $(\Sigma_{d,\Gamma}^*, \sqsubseteq_{p,\Gamma})$.

4.2.1. Controlled Sequences. We employ to this end the framework and results of [SS11]. The first observation is that bad sequences over $(\Sigma_d^*, \sqsubseteq_p)$ can be of arbitrary length: for every $N > 0$, the sequence

$$1, 0^N, 0^{N-1}, \dots, 0$$

is indeed a bad sequence of length $N + 1$ over $(\Sigma_1^*, \sqsubseteq_p)$. Thankfully, what we are looking for are not general bounds over all the bad sequences, but over the kind of sequences that arise in the algorithms of Theorem 2.13: in particular, the bounds can take into account how fast the lengths of the strings in the sequence can grow. Define a *normed* wqo as a wqo (A, \leq_A) further equipped with a norm $|\cdot|_A: A \rightarrow \mathbb{N}$. As a sanity condition, we ask for

$$A_{\leq n} \stackrel{\text{def}}{=} \{x \in A \mid |x|_A \leq n\} \quad (4.4)$$

¹Note that, at such high complexities, the usual distinctions between deterministic vs. nondeterministic, or time-bounded vs. space-bounded computations become irrelevant.

to be finite for each n .

The wqos introduced in Section 2.2 can be normed for instance by

$$\begin{aligned} |a|_{\Sigma} &\stackrel{\text{def}}{=} 0, & |\langle x, i \rangle|_{A_1+A_2} &\stackrel{\text{def}}{=} |x|_{A_i}, \\ |\langle x, y \rangle|_{A_1 \times A_2} &\stackrel{\text{def}}{=} \max(|x|_{A_1}, |y|_{A_2}), & |x_1 \cdots x_{\ell}|_{A_1^*} &\stackrel{\text{def}}{=} \max_{1 \leq i \leq \ell}(\ell, |x_i|_{A_1}), \end{aligned}$$

where $(\Sigma, =)$ denotes a finite set with equality. For $(\Sigma_{d,\Gamma}^*, \sqsubseteq_{p,\Gamma})$, we choose similarly

$$|(a_1, w_1) \cdots (a_{\ell}, w_{\ell})|_{\Sigma_{d,\Gamma}^*} \stackrel{\text{def}}{=} \max_{1 \leq i \leq \ell}(\ell, |w_i|_{\Gamma}), \quad (4.5)$$

where we assume (Γ, \leq_{Γ}) to be normed by $|\cdot|_{\Gamma}$. By the definition above, this simplifies to

$$|(a_1, w_1) \cdots (a_{\ell}, w_{\ell})|_{\Sigma_{d,\Gamma}^*} = \ell \quad (4.6)$$

when Γ is finite.

Let $g: \mathbb{N} \rightarrow \mathbb{N}$ be a strictly monotone function (hereafter called a *control function*), and n be a non-negative integer. A sequence x_0, x_1, \dots over $(A, \leq_A, |\cdot|_A)$ is (g, n) -controlled if, for all i , $|x_i|_A \leq g^i(n)$, the i th iterate of g on n . Note in particular that this entails $|x_0|_A \leq n$. Given an algorithm that relies on (A, \leq_A) being a wqo for its termination, *i.e.*, on the fact that bad sequences over (A, \leq_A) are finite, the intuition is that g should bound how fast the norm of the elements in our bad sequences can grow, and n should bound the norm of the initial element. As shown in [SS11], for a given g and n , bad (g, n) -controlled sequences over $(A, \leq_A, |\cdot|_A)$ have a maximal length denoted $L_{A,g}(n)$.

4.2.2. Normed Reflections. The *Length Function Theorem* in [SS11] provides suitable sub-recursive upper bounds on the function $L_{A,g}$ when A is constructed using the *elementary* wqo algebra that allows disjoint unions, Cartesian products and Kleene star to be used over finite sets. We are going to exploit these bounds together with the order reflection employed in the proof of Theorem 3.6 to obtain a bound on $L_{\Sigma_{d,\Gamma}^*,g}$.

A reflection $r: A \rightarrow B$ between two normed wqos $(A, \leq_A, |\cdot|_A)$ and $(B, \leq_B, |\cdot|_B)$ is *normed* if $|r(x)|_B \leq |x|_A$ for all x in A . We write “ $A \hookrightarrow B$ ” if there exists such a normed reflection from A to B . Observe that, if x_0, x_1, \dots is a (g, n) -controlled bad sequence over A , then $r(x_0), r(x_1), \dots$ is also a (g, n) -controlled bad sequence, this time over B . Thus $A \hookrightarrow B$ implies $L_{A,g} \leq L_{B,g}$, *i.e.*, $L_{A,g}(m) \leq L_{B,g}(m)$ for all $m \in \mathbb{N}$.

One can check, using induction on d , that the reflection $r: \Sigma_{d,\Gamma}^* \rightarrow \Theta_{d,\Gamma}$ used in the proof of Theorem 3.6 is normed. If x is in $\Sigma_{d-1,\Gamma}^*$, then $r(x) = x$ itself with $|x|_{\Sigma_{d,\Gamma}^*} = |r(x)|_{\Theta_{d,\Gamma}}$. Otherwise, let x be factorized as in (3.7); $r(x)$ is given in (3.9). Then on the one hand

$$|x|_{\Sigma_{d,\Gamma}^*} = \max_{0 \leq j < k \leq m} \left(m + \sum_{i=0}^m |x_i|, |v_k|_{\Gamma}, |x_j|_{\Sigma_{d-1,\Gamma}^*} \right), \quad (4.7)$$

while on the other hand

$$|r(x)|_{\Theta_{d,\Gamma}} = \max \left(|x_0|_{\Sigma_{d-1,\Gamma}^*}, \max_{1 \leq i \leq m-1} (m-1, |v_i|_{\Gamma}, |x_i|_{\Sigma_{d-1,\Gamma}^*}), |w_m|_{\Gamma}, |x_m|_{\Sigma_{d-1,\Gamma}^*} \right), \quad (4.8)$$

which indeed satisfy

$$|x|_{\Sigma_{d,\Gamma}^*} \geq |r(x)|_{\Theta_{d,\Gamma}}, \quad (4.9)$$

thanks to the ind. hyp. Therefore $\Sigma_{d,\Gamma}^* \hookrightarrow \Theta_{d,\Gamma}$ for all d, Γ . Define now

$$\Theta'_{-1,\Gamma} \stackrel{\text{def}}{=} \mathbf{1} \quad \Theta'_{d,\Gamma} \stackrel{\text{def}}{=} \Theta'_{d-1,\Gamma} + \Theta'_{d-1,\Gamma} \times (\Gamma \times \Theta'_{d-1,\Gamma})^* \times \Gamma \times \Theta'_{d-1,\Gamma}, \quad (4.10)$$

where $\mathbf{1}$ denotes the singleton set. Since \hookrightarrow is a precongruence for the elementary algebraic operations [SS11, Proposition 3.5], we deduce that $\Sigma_{d,\Gamma}^* \hookrightarrow \Theta'_{d,\Gamma}$ and thus

$$L_{\Sigma_{d,\Gamma}^*,g} \leq L_{\Theta'_{d,\Gamma},g} \quad (4.11)$$

for all normed wqos Γ , all d , and all control functions g . Assuming $(\Gamma, \leq_\Gamma, |\cdot|_\Gamma)$ to be elementary, then each $\Theta'_{d,\Gamma}$ is also elementary, *i.e.* we are going to be able to apply the Length Function Theorem to it and derive an upper bound for $L_{\Theta'_{d,\Gamma},g}$, and thereby for $L_{\Sigma_{d,\Gamma}^*,g}$.

4.2.3. Maximal Order Types. The version of the Length Function Theorem we wish to apply requires the computation of the *maximal order type* of $\Theta'_{d,\Gamma}$. This is a measure of the complexity of a wqo (A, \leq_A) defined in [dJP77] as the maximal order type of its linearizations: a *linearization* \prec of \leq_A is a total linear ordering over A that contains $\leq_A \setminus \geq_A$ as a subrelation. Any such linearization of a wqo is well-founded and thus isomorphic to an ordinal, called its order type, and the maximal order type of (A, \leq_A) is therefore the maximal such ordinal.

De Jongh and Parikh [dJP77] provide formulæ to compute the maximal order types of elementary wqos based on their algebraic decompositions as disjoint sums, Cartesian products, and Kleene star—using respectively the sum ordering, the product ordering, and the subword embedding ordering—:

$$\begin{aligned} o(A + B) &= o(A) \oplus o(B), \\ o(A \times B) &= o(A) \otimes o(B), \\ o(A^*) &= \begin{cases} \omega^{\omega^{o(A)}-1} & \text{if } A \text{ is finite,} \\ \omega^{\omega^{o(A)}} & \text{otherwise.} \end{cases} \end{aligned}$$

Here, the \oplus and \otimes operations are the *natural sum* and *natural product* on ordinals, defined for ordinals in CNF in ε_0 by

$$\sum_{i=1}^m \omega^{\beta_i} \oplus \sum_{j=1}^n \omega^{\beta'_j} \stackrel{\text{def}}{=} \sum_{k=1}^{m+n} \omega^{\gamma_k}, \quad \sum_{i=1}^m \omega^{\beta_i} \otimes \sum_{j=1}^n \omega^{\beta'_j} \stackrel{\text{def}}{=} \bigoplus_{i=1}^m \bigoplus_{j=1}^n \omega^{\beta_i \oplus \beta'_j}, \quad (4.12)$$

where $\gamma_1 \geq \dots \geq \gamma_{m+n}$ is a reordering of $\beta_1, \dots, \beta_m, \beta'_1, \dots, \beta'_n$.

Schütte and Simpson [SS85] compute the exact maximal order type of a wqo related to $(\Sigma_{d,\Gamma}^*, \sqsubseteq_p)$. Here we are content with the maximal order type of $\Theta'_{d,\Gamma}$ (which also provides an upper bound on $o(\Sigma_{d,\Gamma}^*)$). By (4.11) and [SS11, Proposition 5.2], we obtain:²

Proposition 4.1 (Length Function Theorem for $\Sigma_{d,\Gamma}^*$). *Let $d \in \mathbb{N}$ and assume Γ is an elementary wqo and g is a control function. Then there exists a polynomial p independent of d, Γ, g such that $L_{\Sigma_{d,\Gamma}^*,g} \leq (p \circ g)^{o(\Theta'_{d,\Gamma})}$.*

²To be precise, [SS11, Proposition 5.2] only provides bounds for *exponential* wqos—where there are no nested applications of the Kleene star operation—but it can be generalized to elementary wqos.

4.2.4. *Finite Alphabets and Successor Control.* Proposition 4.1 is more general than useful for deriving upper bounds on PCSs verification. Even if we use a generalized priority alphabet in our PCSs, by allowing read and write rules to manipulate pairs in $\Sigma_{d,\Gamma}$ instead of only priorities in Σ_d , we can safely assume that the underlying alphabet Γ is finite and part of the input, with maximal order type $o(\Gamma) = |\Gamma|$. Similarly, the successor function $H(x) \stackrel{\text{def}}{=} x + 1$ can be chosen for the control function g , thus $p \circ g$ in Proposition 4.1 is simply a polynomial.

We can furthermore simplify the ordinal index in Proposition 4.1. First note that

$$o(\Theta'_{d,\Gamma}) < (\Omega_{2(d+1)+1})_{|\Gamma|} \quad (4.13)$$

for all d in \mathbb{N} , where $(\Omega_{2(d+1)+1})_{|\Gamma|} = \omega^{\dots\omega^{|\Gamma|}} \} 2(d+1)$ stacked ω 's. Second, an ordinal term α in CNF can be written as $\alpha = \omega^{\alpha_1} \cdot c_1 + \dots + \omega^{\alpha_m} \cdot c_m$ for $\alpha > \alpha_1 > \dots > \alpha_m$ and $0 < c_1, \dots, c_m < \omega$. We define then inductively its *maximum coefficient* $N(\alpha)$ as $\max(N(\alpha_1), \dots, N(\alpha_m), c_1, \dots, c_m)$. Observe that for all d in \mathbb{N} ,

$$N(o(\Theta'_{d,\Gamma})) \leq |\Gamma|. \quad (4.14)$$

The following simplified statement then holds:

Corollary 4.2. *Let $d \in \mathbb{N}$ and Γ be a finite non-empty alphabet. Then there exists a polynomial h independent of d, Γ such that $L_{\Sigma_{d,\Gamma}^*, H}(n) \leq h^{\Omega_{2(d+1)+1}}(n)$ for all $n \geq |\Gamma|$.*

Proof. By Proposition 4.1 it suffices to show that $h^{o(\Theta'_{d,\Gamma})}(n) \leq h^{\Omega_{2(d+1)+1}}(n)$ for all $n \geq |\Gamma| > 0$. We show instead $h^{o(\Theta'_{d,\Gamma})}(n) \leq h^{(\Omega_{2(d+1)+1})_{|\Gamma|}}(n)$ since it allows to conclude. By monotonicity of the Hardy functions in the ordinal index for the *pointwise ordering*—see [FW98, Theorem 2.21.2] or [SS12, Lemma A.10]—it suffices to show that $o(\Theta'_{d,\Gamma}) \preceq_{|\Gamma|} (\Omega_{2(d+1)+1})_{|\Gamma|}$ using the notation of [SS12], which is entailed by (4.13–4.14) and [SS12, Lemma A.5]. \square

4.3. **Complexity Upper Bounds.** Now that we are armed with a Length Function Theorem for $(\Sigma_{d,\Gamma}^*, \sqsubseteq_{p,\Gamma})$, we can prove an upper bound for PCS verification:

Theorem 4.3 (Complexity of PCS Verification). *Reachability and Inevitability of PCSs are in $\mathbf{F}_{\varepsilon_0}$.*

Let us explain the steps towards an upper bound for Termination in some detail; the results for Reachability and Inevitability are similar but more involved—see [SS12, SS13] for generic complexity arguments for WSTSs.

A Finite Witness. Observe that, if an execution $C_0 \rightarrow_{\#} C_1 \rightarrow_{\#} C_2 \rightarrow_{\#} \dots$ of the transition system $\mathcal{S}_{\#}$ satisfies $C_i \leq_{\#} C_j$ for some indices $i < j$, then because $\mathcal{S}_{\#}$ is a WSTS, we can simulate the steps performed in this sequence after C_i but starting from C_j and build an infinite run. Conversely, if the system does not terminate, *i.e.* if there is an infinite execution $C_0 \rightarrow_{\#} C_1 \rightarrow_{\#} C_2 \rightarrow_{\#} \dots$, then because of the wqo we will eventually find $i < j$ such that $C_i \leq_{\#} C_j$. Therefore, the system is non-terminating if and only if there is a finite witness of the form $C_0 \xrightarrow{*}_{\#} C_i \xrightarrow{\dagger}_{\#} C_j$ with $C_i \leq_{\#} C_j$.

Controlled Witnesses. Another observation is that the size of successive configurations cannot grow arbitrarily along runs; in fact, the length of the channels contents can only grow by one symbol at a time using a write transition. This means that if we define $|C = (q, x_1, \dots, x_m)| = \sum_{j=1}^m |x_j|$, then in an execution $C_0 \rightarrow_{\#} C_1 \rightarrow_{\#} C_2 \rightarrow_{\#} \dots$, $|C_i| \leq |C_0| + i = H^i(|C_0|)$, *i.e.* any execution is *controlled* by the successor function H .

Applying the Length Function Theorem. Corollary 4.2 yields an $h^{\Omega_{2(d+1)+1}}(|C_0| + |\Gamma|)$ upper bound on the length of bad $(H, |C_0|)$ -controlled sequences over $(\Sigma_{d,\Gamma}^*, \sqsubseteq_{p,\Gamma})$ for some polynomial h . It can be lifted to bound the maximal length of a witness in $\mathcal{S}_{\#}$, when considering instead the ordinal $o_S \stackrel{\text{def}}{=} (\Omega_{2(d+1)+1})^m \cdot |Q|$. Setting $|S| = |\Delta| + |Q| + d + m + |\Gamma|$, we see that this length is less than $H^{\varepsilon_0}(p(|S| + |C_0|)) < F_{\varepsilon_0}(p(|S| + |C_0|))$ for some fixed ordinal-recursive function p .

A Combinatorial Algorithm. Since the functions $(h^\alpha)_\alpha$ are elementary constructive [Sch13, Theorem 5.1], the above discussion yields a non-deterministic algorithm in $\mathbf{F}_{\varepsilon_0}$ for Termination: compute $L = h^{o_S}(|C_0| + |\Gamma|)$ and look for an execution of length $L + 1$ in $\mathcal{S}_{\#}$. If one exists, it is necessarily a witness for nontermination; otherwise, the system is guaranteed to terminate from C_0 .

We call this a *combinatorial* algorithm, as it relies on the combinatorial analysis provided by the Length Function Theorem to derive an upper bound on the size of a finite witness for the property at hand—here Termination, but the same kind of techniques can be used for Reachability and Inevitability.

5. HARDY COMPUTATIONS BY PCSs

In this section we show how PCSs can weakly compute the Hardy functions H^α and their inverses for all ordinals α below Ω , which is the key ingredient for Theorem 5.8 below stating our hardness result. For this, we develop in Section 5.1 encodings $s(\alpha) \in \Sigma_d^*$ for ordinals $\alpha \in \Omega_d$ and show how PCSs can compute with these codes, *e.g.* build the code for λ_n from the code of a limit λ . This is used in Section 5.3 to design PCSs that *weakly compute* H^α and $(H^\alpha)^{-1}$ in the sense of Definition 5.5 below.

5.1. Encoding Ordinals. Our encoding of ordinal terms as strings in Σ_d^* employs strings of a particular form. For $0 \leq a \leq d$, we use the following equation to define the language $C_a \subseteq \Sigma_d^*$ of *codes*:

$$C_a \stackrel{\text{def}}{=} \varepsilon + C_a C_{a-1} a, \quad C_{-1} \stackrel{\text{def}}{=} \varepsilon. \quad (5.1)$$

Let $C = C_{-1} + C_0 + \dots + C_d$. Each C_a (and then C itself) is a regular language, with $C_a = (C_{a-1} a)^*$; for instance, $C_0 = 0^*$.

5.1.1. *Decompositions.* A code x is either the empty word ε , or belongs to a unique C_a . If $x \in C_a$ is not empty, it has a unique factorization $x = yza$ according to (5.1) with $y \in C_a$ and $z \in C_{a-1}$. Recall that $h(x)$ denotes the height function, thus a non-empty $x = a_1 \cdots a_\ell$ is a code if and only if $a_\ell = h(x)$ and $a_{i+1} - a_i \leq 1$ for all $i < \ell$ (we say that x has no jumps: priorities only increase smoothly along codes, but they can decrease sharply). For instance, 02 is not a code (it has a jump), but 001122 and 01223400123334 are codes.

The factor $z \in C_{a-1}$ in $x = yza$ can be developed further, as long as $z \neq \varepsilon$: a non-empty code $x \in C_d$ has a unique factorization as $x = y_d y_{d-1} \cdots y_a a^d$ with $y_i \in C_i$ for $i = a, \dots, d$, and where for $0 \leq a \leq b$, we write a^b for the staircase word $a(a+1) \cdots (b-1)b$, letting $a^b = \varepsilon$ when $a > b$. We call this the *decomposition* of x . Note that the value of a is obtained by looking for the maximal suffix of x that is a staircase word. For example, $x = 23312340121234 \in C_4$ is a code and decomposes as

$$x = \overbrace{2331234}^{y_4} \overbrace{\varepsilon}^{y_3} \overbrace{012}^{y_2} \overbrace{\varepsilon}^{y_1} \overbrace{1234}^{1^4}.$$

5.1.2. *Ordinal Encoding.* With a code $x \in C$, we associate an ordinal term $\eta(x)$ given by

$$\eta(\varepsilon) \stackrel{\text{def}}{=} 0, \quad \eta(yza) \stackrel{\text{def}}{=} \eta(y) + \omega^{\eta(z)}, \quad (5.2)$$

where $x = yza$ is the factorization according to (5.1) of $x \in C_a \setminus \{\varepsilon\}$. For example, $\eta(a) = \omega^0 = 1$ for all $a \in \Sigma_d$, $\eta(012) = \eta(234) = \omega^\omega$, and more generally $\eta(a^b) = \Omega_{b-a}$. One sees that $\eta(x) < \Omega_{a+1}$ when $x \in C_a$.

This *decoding* function $\eta: C \rightarrow \Omega_{d+1}$ is onto (or surjective) but it is not bijective. However, it is a bijection between C_a and Ω_{a+1} for any $a \leq d$. Its converse is the level- a *encoding* function $s_a: \Omega_{a+1} \rightarrow C_a$, defined with

$$s_a\left(\sum_{i=1}^p \gamma_i\right) \stackrel{\text{def}}{=} s_a(\gamma_1) \cdots s_a(\gamma_p), \quad s_a(\omega^\alpha) \stackrel{\text{def}}{=} s_{a-1}(\alpha) a. \quad (5.3)$$

Thus $s_a(0) = \varepsilon$ and, for example,

$$\begin{aligned} s_5(1) &= 5, & s_5(3) &= 555, & s_5(\omega) &= 45, \\ s_5(\omega^3) &= 4445, & s_5(\omega^\omega) &= 345, & s_5(\omega^{\omega^\omega}) &= 2345, \\ s_5(\omega^3 + \omega^2) &= 4445445, & s_5(\omega \cdot 3) &= 454545. \end{aligned}$$

We may omit the subscript when $a = d$, *e.g.* writing $s(1) = d$.

5.1.3. *Successors and Limits.* Let $x = y_d y_{d-1} \cdots y_a a^d$ be the decomposition of $x \in C_d \setminus \{\varepsilon\}$. By (5.2), x encodes a successor ordinal $\eta(x) = \beta + 1$ if and only if $a = d$, *i.e.*, if x ends with two d 's (or has length 1). Since then $\beta = \eta(y_d \cdots y_a)$, one obtains the “predecessor of x ” by removing the final d .

If $a < d$, x encodes a limit λ . Combining (4.1) and (5.2), one obtains the encoding $(x)_n$ of λ_n with

$$(x)_n = y_d y_{d-1} \cdots y_{a+1} (y_a (a+1))^n (a+2)^d. \quad (5.4)$$

For instance, with $d = 5$, decomposing $x = 333345 = s(\omega^{\omega^4})$ gives $a = 3$, $x = y_5 y_4 y_3 3^5$, with $y_3 = 333$ and $y_5 = y_4 = \varepsilon$. Then $(x)_n = (3334)^n 5$, agreeing with, *e.g.* $s(\omega^{\omega^{3 \cdot 2}}) = 333433345$.

5.2. Robustness. A crucial property of our ordinal encoding is *robustness*, i.e. that $x \sqsubseteq_p x'$ should reflect the corresponding relation $H^{\eta(x)}(n) \leq H^{\eta(x')}(n)$ on Hardy computations.

Proposition 5.1 (Robustness). *Let $a \geq 0$ and $x \sqsubseteq_p x'$ be two strings in C_a . Then, $H^{\eta(x)}(n) \leq H^{\eta(x')}(n')$ for all $n \leq n'$ in \mathbb{N} .*

The proof of Proposition 5.1 requires delving in some of the theory of subrecursive functions, and is postponed until Section 5.2.3.

5.2.1. Properties of the Hardy Hierarchy. We first list some useful properties of Hardy computations (see [FW98] or [SS12, App. A] for details). The first fact is that each Hardy function is expansive and monotone in its argument n :

Fact 5.2 (Expansiveness and Monotonicity). *For all α, α' in Ω and $n > 0, m$ in \mathbb{N} ,*

$$n \leq H^\alpha(n), \quad (5.5)$$

$$n \leq m \text{ implies } H^\alpha(n) \leq H^\alpha(m). \quad (5.6)$$

However, the Hardy functions are not monotone in the ordinal parameter: $H^{n+1}(n) = 2n+1 > 2n = H^n(n) = H^\omega(n)$, though $n+1 < \omega$. We will introduce an ordering on ordinal terms in Section 5.2.2 that ensures monotonicity of the Hardy functions.

Another handful fact is that we can decompose Hardy computations:

Fact 5.3. *For all α, γ in Ω , and n in \mathbb{N} ,*

$$H^{\gamma+\alpha}(n) = H^\gamma(H^\alpha(n)). \quad (5.7)$$

Note that (5.7) holds for all ordinal terms, and not only for those α, γ such that $\gamma + \alpha$ is in CNF—this is a virtue of working with terms rather than set-theoretic ordinals.

5.2.2. Ordinal Embedding. We introduce a partial ordering \sqsubseteq_o on ordinal terms, called *embedding*, and which corresponds to a strict tree embedding on the structure of ordinal terms. Formally, it is defined inductively by $\alpha \sqsubseteq_o \beta \stackrel{\text{def}}{\iff} \alpha = \omega^{\alpha_1} + \dots + \omega^{\alpha_p}, \beta = \omega^{\beta_1} + \dots + \omega^{\beta_m}$, and there exist $1 \leq i_1 < i_2 < \dots < i_p \leq m$ such that $\alpha_1 \sqsubseteq_o \beta_{i_1} \wedge \dots \wedge \alpha_p \sqsubseteq_o \beta_{i_p}$. Note that $0 \sqsubseteq_o \alpha$ for all α , that $1 \sqsubseteq_o \alpha$ for all $\alpha > 0$. In general, $\alpha \not\sqsubseteq_o \omega^\alpha$ and $\lambda_n \not\sqsubseteq_o \lambda$. Ordinal embedding is congruent for addition and ω -exponentiation of terms:

$$\alpha \sqsubseteq_o \alpha' \text{ and } \beta \sqsubseteq_o \beta' \text{ imply } \alpha + \beta \sqsubseteq_o \alpha' + \beta', \quad (5.8)$$

$$\alpha \sqsubseteq_o \alpha' \text{ implies } \omega^\alpha \sqsubseteq_o \omega^{\alpha'}, \quad (5.9)$$

and could in fact be defined alternatively by the axiom $0 \sqsubseteq_o \alpha$ and the two deduction rules (5.8) and (5.9).

We list a few useful consequences of the definition of \sqsubseteq_o :

$$\alpha \sqsubseteq_o \gamma + \omega^\beta \text{ implies } \alpha \sqsubseteq_o \gamma, \text{ or } \alpha = \gamma' + \omega^{\beta'} \text{ with } \gamma' \sqsubseteq_o \gamma \text{ and } \beta' \sqsubseteq_o \beta, \quad (5.10)$$

$$n \leq m \text{ implies } \lambda_n \sqsubseteq_o \lambda_m, \quad (5.11)$$

$$\alpha \sqsubseteq_o \lambda \text{ implies } \alpha \sqsubseteq_o \lambda_n, \text{ or } \alpha \text{ is a limit and } \alpha_n \sqsubseteq_o \lambda_n. \quad (5.12)$$

Proof of (5.10). Intuitively, there are two cases when we consider $\alpha \sqsubseteq_o \alpha' = \gamma + \omega^\beta$: either the ω^β summand of α' is in the range of the embedding or not. If it is not, then already $\alpha \sqsubseteq_o \gamma$. If it is, then α must be some $\gamma' + \omega^{\beta'}$ and $\omega^{\beta'} \sqsubseteq_o \omega^\beta$, which implies in turn $\beta' \sqsubseteq_o \beta$. \square

Proof of (5.11). By induction on λ : indeed if $\lambda = \gamma + \omega^{\beta+1}$ then $\lambda_m = \gamma + \omega^\beta \cdot m$, which is $\lambda_n + \omega^\beta \cdot (m - n)$. If $\lambda = \gamma + \omega^{\lambda'}$, the ind. hyp. gives $\lambda'_n \sqsubseteq_o \lambda'_m$, hence $\lambda_n = \gamma + \omega^{\lambda'_n} \sqsubseteq_o \gamma + \omega^{\lambda'_m} = \lambda_m$. \square

Proof of (5.12). By induction on λ . We can write λ as some $\gamma + \omega^\beta$ with $\beta > 0$ so that $\lambda_n = \gamma + (\omega^\beta)_n$. If $\alpha \sqsubseteq_o \gamma$, then $\alpha \sqsubseteq_o \lambda_n$ trivially. If $\alpha = \gamma' + 1$ is a successor, $1 \sqsubseteq_o (\omega^\beta)_n$ and again $\alpha \sqsubseteq_o \lambda_n$. There remains the case where $\alpha = \gamma' + \omega^{\beta'}$ is a limit (*i.e.* $\beta' > 0$) with $\gamma' \sqsubseteq_o \gamma$ and $\beta' \sqsubseteq_o \beta$. If β is a limit, then by ind. hyp. either $\beta' \sqsubseteq_o \beta_n$ and hence $\alpha \sqsubseteq_o \lambda_n$, or β' is a limit and $\beta'_n \sqsubseteq_o \beta_n$, hence $\alpha_n \sqsubseteq_o \lambda_n$. Finally, if $\beta = \delta + 1$ is a successor, then either $\beta' \sqsubseteq_o \delta$ so that $\alpha \sqsubseteq_o \gamma + \omega^\delta \sqsubseteq_o \gamma + \omega^\delta \cdot n = \lambda_n$, otherwise by (5.10), β' is a successor $\delta' + 1$ with $\delta' \sqsubseteq_o \delta$, and then $(\omega^{\beta'})_n = \omega^{\delta'} \cdot n \sqsubseteq_o \omega^\delta \cdot n = (\omega^\beta)_n$, hence $\alpha_n \sqsubseteq_o \lambda_n$. \square

Proposition 5.4 (Monotonicity). *For all α, α' in Ω and n in \mathbb{N} ,*

$$\alpha \sqsubseteq_o \alpha' \text{ implies } H^\alpha(n) \leq H^{\alpha'}(n).$$

Proof. Let us proceed by induction on a proof of $\alpha \sqsubseteq_o \alpha'$, based on the deduction rules (5.8) and (5.9). For the base case, $0 \sqsubseteq_o \alpha'$ implies $H^0(n) = n \leq H^{\alpha'}(n)$ by expansiveness.

For the inductive step with (5.8), if $\alpha \sqsubseteq_o \alpha'$ and $\beta \sqsubseteq_o \beta'$, then

$$\begin{aligned} H^{\alpha+\beta}(n) &= H^\alpha(H^\beta(n)) && \text{by (5.7)} \\ &\leq H^\alpha(H^{\beta'}(n)) && \text{by ind. hyp. and (5.6)} \\ &\leq H^{\alpha'}(H^{\beta'}(n)) && \text{by ind. hyp.} \\ &= H^{\alpha'+\beta'}(n). && \text{by (5.7)} \end{aligned}$$

For the inductive step with (5.9), if $\alpha \sqsubseteq_o \alpha'$, then we show $H^{\omega^\alpha}(n) \leq H^{\omega^{\alpha'}}(n)$ by induction on α' :

- If $\alpha' = 0$, then $\alpha = 0$ and we are done.
- If $\alpha' = \beta' + 1$ is a successor, then by (5.10) either $\alpha \sqsubseteq_o \beta'$, or $\alpha = \beta + 1$ with $\beta \sqsubseteq_o \beta'$. In the first case, $H^{\omega^\alpha}(n) \leq H^{\omega^{\beta'}}(n) \leq H^{\omega^{\beta'}}(H(n)) = H^{\omega^{\alpha'}}(n)$ by ind. hyp. and expansiveness. In the second case, we see by induction on $i \in \mathbb{N}$ that

$$\left(H^{\omega^\beta}\right)^i(n) \leq \left(H^{\omega^{\beta'}}\right)^i(n) \tag{5.13}$$

for all i and n thanks to the ind. hyp. Thus

$$H^{\omega^{\beta+1}}(n) = \left(H^{\omega^\beta}\right)^n(n) \leq \left(H^{\omega^{\beta'}}\right)^n(n) = H^{\omega^{\beta'+1}}(n)$$

for all n , and we are done.

- If $\alpha' = \lambda'$ is a limit, then by (5.12) either $\alpha \sqsubseteq_o \lambda'_n$ or α is a limit λ and $\lambda_n \sqsubseteq_o \lambda'_n$. In the first case $H^{\omega^\alpha}(n) \leq H^{\omega^{\lambda'_n}}(n)$ by ind. hyp.; in the second case $H^{\omega^\lambda}(n) = H^{\omega^{\lambda'_n}}(n) \leq H^{\omega^{\lambda'_n}}(n) = H^{\omega^{\lambda'}}(n)$ using the ind. hyp. \square

o :	334545\$	the <i>ordinal</i> term $\omega^{\omega^2} + \omega^\omega$
c :	0000\$	the <i>counter</i> value 4
t :	\$	the <i>temporary</i> storage

FIGURE 2. Channels for Hardy computations.

5.2.3. *Robustness.* We are now in position to prove Proposition 5.1:

Proof of Proposition 5.1. We prove that $\eta(x) \sqsubseteq_o \eta(x')$ by induction on x and conclude using Proposition 5.4 and Eq. (5.6). If $x = \varepsilon$, $\eta(x) = 0 \sqsubseteq_o \eta(x')$. Otherwise we can decompose x as yza according to (5.1) with $y \in C_a$ and $z \in C_{a-1}$. By (3.5), $x' = y'z'a$ with $y \sqsubseteq_p y'$ and $za \sqsubseteq_p z'a$. Observe that y' and z' are in C_a , and writing $z'a = z'_1 a \cdots z'_m a$ for the canonical decomposition of z' —where necessarily each z'_j is in C_{a-1} —, then $z \sqsubseteq_p z'_1$ as there is no other way of disposing of the other occurrences of a in z' .

By ind. hyp., $\eta(y) \sqsubseteq_o \eta(y')$ and $\eta(z) \sqsubseteq_o \eta(z'_1)$. Then, because $\eta(x) = \eta(y) + \omega^{\eta(z)}$ and $\eta(x') = \eta(y') + \omega^{\eta(z'_1)} + \cdots + \omega^{\eta(z'_m)}$, we see by (5.8) and (5.9) that $\eta(x) \sqsubseteq_o \eta(x')$. \square

5.3. **Robust Hardy Computations in PCSs.** Our goal is to implement in a PCS the *canonical* Hardy steps, denoted with \xrightarrow{H} , and specified by the following two rewrite rules on pairs in $\Omega \times \mathbb{N}$:

$$(\alpha + 1, n) \xrightarrow{H} (\alpha, n + 1) \quad \text{for successors,} \quad (5.14)$$

$$(\lambda, n) \xrightarrow{H} (\lambda_n, n) \quad \text{for limits.} \quad (5.15)$$

A *Hardy computation* for $H^\alpha(n)$ is a sequence of rewrites $(\alpha, n) = (\alpha_0, n_0) \xrightarrow{H} (\alpha_1, n_1) \xrightarrow{H} \cdots \xrightarrow{H} (\alpha_\ell, n_\ell)$. Note that (5.14–5.15) provide a rewriting view of the definition of Hardy functions in (4.2). Thus $H^{\alpha_i}(n_i)$ remains constant throughout the computation, and if $\alpha_\ell = 0$ then $n_\ell = H^\alpha(n)$, in which case we call the computation *complete*.

We do not implement canonical Hardy steps as PCSs, but construct instead *weak* computers, which might return lower values. Our PCSs for weak Hardy computations use three channels (see Figure 2), storing (codes for) a pair α, n on channels **o** (for “ordinal”) and **c** (for “counter”), and employ an extra channel, **t**, for “temporary” storage. Instead of Σ_d , we use Σ_{d+1} with $d+1$ used as a position marker and written \$ for clarity: each channel always contains a single occurrence of \$.

Definition 5.5. A *weak Hardy computer* for Ω_{d+1} is a $(d+1)$ -PCS S with channels $\mathbf{Ch} = \{\mathbf{o}, \mathbf{c}, \mathbf{t}\}$ and two distinguished states p_{beg} and p_{end} such that:

$$\text{if } (p_{\text{beg}}, x\$, y\$, z\$, \$) \xrightarrow{*}_{\#} (p_{\text{end}}, u, v, w) \quad \text{then } x \in C_d, y \in 0^+, z = \varepsilon, u, v, w \in \Sigma_d^*\$, \quad (\text{safety})$$

$$\text{if } (p_{\text{beg}}, s(\alpha)\$, 0^n\$, \$) \xrightarrow{*}_{\#} (p_{\text{end}}, s(\beta)\$, 0^m\$, \$) \quad \text{then } H^\alpha(n) \geq H^\beta(m). \quad (\text{robustness})$$

Furthermore S is *complete* if for any $\alpha < \Omega_{d+1}$ and $n > 0$,

$$(p_{\text{beg}}, s(\alpha)\$, 0^n\$, \$) \xrightarrow{*}_{\#} (p_{\text{end}}, \$, 0^m\$, \$) \quad \text{where } m \stackrel{\text{def}}{=} H^\alpha(n) \quad (\text{complete})$$

and it is *inv-complete* if

$$(p_{\text{beg}}, \$, 0^m \$, \$) \xrightarrow{*}_{\#} (p_{\text{end}}, s(\alpha) \$, 0^n \$, \$). \quad (\text{inv-complete})$$

Lemma 5.6 (PCSs weakly compute Hardy functions). *For every $d \in \mathbb{N}$, there exists a weak Hardy computer S_d for Ω_{d+1} that is complete, and a weak S_d^{-1} that is inv-complete. Furthermore S_d and S_d^{-1} can be generated uniformly in $\text{DSpace}(\log d)$.*

We design a complete weak Hardy computer for Lemma 5.6 by assembling several components. The weak Hardy computer S_d is actually composed of two components $S_{d,+1}$ and $S_{d,\lambda}$ in charge respectively of applying the successor (5.14) and limit (5.15) steps, and similarly S_d^{-1} is composed of two components $S_{d,+1}^{-1}$ and $S_{d,\lambda}^{-1}$ in charge of reversing those steps.

5.3.1. *Successor Steps.* We start with “canonical successor steps”, as *per* (5.14). They are implemented by $S_{d,+1}$, the PCS depicted in Figure 3. When working on codes, replacing $s(\alpha + 1)$ by $s(\alpha)$ simply means removing the final d (see Section 5.1.3), but when the strings are in fifo channels this requires reading the whole contents of a channel and writing them back, relying on the $\$$ end-marker.




FIGURE 3. $S_{d,+1}$, a PCS for Hardy steps $(\alpha + 1, n) \xrightarrow{H} (\alpha, n + 1)$.

Remark 5.7 (Notational/graphical conventions). The label edge “ $q \xrightarrow{c ?! 0} q$ ” in Figure 3, with $c ?! 0$ as label, is shorthand notation for “ $q \xrightarrow{c ?! 0} o \xrightarrow{c ! 0} q$ ”, letting the intermediary state remain implicit. We also use meta-rules like “ $p \xrightarrow{o ?! x \in C_d} o$ ” above to denote a subsystem tasked with reading and writing back a string x over o while checking that it belongs to C_d ; since C_d is a regular language, such subsystems are directly obtained from DFAs for C_d . \square

We first analyze the behavior of $S_{d,+1}$ when superseding of low-priority messages does not occur, *i.e.*, we first consider its “reliable” semantics. In this case, starting $S_{d,+1}$ in state p performs the step given in (5.14) for successor ordinals. More precisely, $S_{d,+1}$ guarantees

$$(p, s(\alpha + 1) \$, 0^n \$, \$) \xrightarrow{*}_{\text{rel}} (r, u, v, w) \text{ iff } u = s(\alpha) \$ \wedge v = 0^{n+1} \$ \wedge w = \$. \quad (5.16)$$

Note that (5.16) refers to “ $\xrightarrow{*}_{\text{rel}}$ ”, with no superseding.

Observe that $S_{d,+1}$ has to non-deterministically guess where the end of $s(\alpha)$ occurs before reading $d \$$ in channel o , and will deadlock if it guesses incorrectly. We often rely on this kind of non-deterministic programming to reduce the size of the PCSs we build. Finally, we observe that if x does not end with dd (and is not just d), *i.e.*, if $\eta(x)$ is not a successor ordinal, then $S_{d,+1}$ will certainly deadlock.

We now consider $S_{d,+1}^{-1}$, the PCS depicted in Figure 4 that implements the inverse canonical steps $(\alpha, n + 1) \xrightarrow{H^{-1}} (\alpha + 1, n)$. Implementing such steps on codes is an easy string-rewriting task since $s(\alpha + 1) = s(\alpha)d$, however our PCS must again read the whole




FIGURE 4. $S_{d,+1}^{-1}$, a PCS for inverse Hardy steps $(\alpha, n+1) \xrightarrow{H^{-1}} (\alpha+1, n)$.

contents of its channels, write them back with only minor modifications while fulfilling the safety requirement of Definition 5.5. When considering the reliable behavior, $S_{d,+1}^{-1}$ guarantees

$$(p, x\$, y\$, \$) \xrightarrow{*}_{\text{rel}} (r, u, v, w) \text{ iff } x \in C_d, \exists n : y = 0^{n+1}, u = s(\eta(x)+1)\$, v = 0^n\$, \text{ and } w = \$. \quad (5.17)$$

Consider now the behavior of $S_{d,+1}$ and $S_{d,+1}^{-1}$ when superseding may occur. Note that a run $(p, x\$, y\$, z\$) \xrightarrow{*}_w (r, \dots)$ from p to r is a *single-pass* run: it reads the whole contents of channels o and c once, and writes some new contents. This feature assumes that we start with a single $\$$ at the end of each channel, as expected by $S_{d,+1}$. For such single-pass runs, the PCS behavior with superseding semantics can be derived from the reliable behavior: for single-pass runs, $C \xrightarrow{*}_w D$ if and only if $C \xrightarrow{*}_{\text{rel}} D' \geq_{\#} D$ for some D' .

Combined with (5.16), the above remark entails robustness for $S_{d,+1}$: there is an execution $(p, s(\alpha+1)\$, 0^n\$, \$) \xrightarrow{*}_w (r, s(\beta)\$, 0^{n'}\$, \$)$ if and only if $s(\beta) \leq_{\#} s(\alpha)$ and $0^{n'}\$ \leq_{\#} 0^{n+1}\$,$ i.e., $n' \leq n+1$. With Proposition 5.1, we deduce $H^{\beta}(n') \leq H^{\alpha}(n)$.

The same reasoning applies to $S_{d,+1}^{-1}$ since this PCS also performs single-pass runs from p to r , hence $(p, s(\alpha)\$, 0^n\$, \$) \xrightarrow{*}_w (r, s(\beta)\$, 0^{n'}\$, \$)$ if and only if $s(\beta) \leq_{\#} s(\alpha+1)$ and $n' \leq n-1$. Thus $H^{\beta}(n') \leq H^{\alpha}(n)$.

5.3.2. Limit Steps. Our next component is $S_{d,\lambda}$, see Figure 5, which implements the canonical Hardy steps for limits from (5.15). The construction follows (5.4): $S_{d,\lambda,a}$ reads (and writes back) the contents of channel o ; guessing non-deterministically the decomposition $y_d \dots y_{a+1} y_a a(a+1)^{\wedge d}$ of $s(\lambda)$, it writes back $y_d \dots y_{a+1}$ and copies y_a on the temporary \mathfrak{t} with $a+1$ appended. Then, a loop around state q_a copies 0^n from and back to c . Every time one 0 is transferred, the whole contents of \mathfrak{t} , initialized with $y_a(a+1)$, is copied to o . When the loop has been visited n times, $S_{d,\lambda,a}$ empties \mathfrak{t} and resumes the transfer of $s(\lambda)$ by copying the final $(a+2)^{\wedge d}$.

For clarity, $S_{d,\lambda,a}$ as given in Figure 5 assumes that a is fixed. The actual $S_{d,\lambda}$ component guesses non-deterministically what is the value of a for the $s(\lambda)$ code on o and gives the control to $S_{d,\lambda,a}$ accordingly.

As far as reliable steps are considered, $S_{d,\lambda}$ guarantees

$$(p, s(\alpha)\$, 0^n\$, \$) \xrightarrow{*}_{\text{rel}} (r, u, v, w) \text{ iff } \alpha \in \text{Lim}, u = s(\alpha_n)\$, v = 0^n\$, \text{ and } w = \$. \quad (5.18)$$

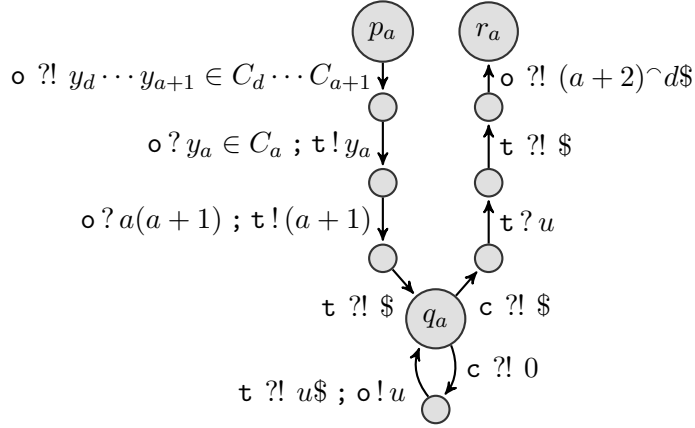


FIGURE 5. $S_{d,\lambda,a}$, a PCS for Hardy steps $(\lambda, n) \xrightarrow{H} (\lambda_n, n)$.

If superseding is allowed, a run $(p_a, s(\alpha)\$, 0^n\$, \$) \xrightarrow{*}_w (r_a, u, v, w)$ has the form

$$\begin{aligned}
(p_a, s(\alpha)\$, 0^n\$, \$) &\xrightarrow{*}_w C_0 = (q_a, (a+2)^\wedge d\$x_0, 0^n\$, z_0\$) \\
&\xrightarrow{*}_w C_1 = (q_a, (a+2)^\wedge d\$x_1, 0^{n-1}\$v_1, z_1\$) \\
&\vdots \\
&\xrightarrow{*}_w C_n = (q_a, (a+2)^\wedge d\$x_n, \$v_n, z_n\$) \\
&\xrightarrow{*}_w (r_a, x'_n\$, v'_n\$, \$)
\end{aligned}$$

where $C_i = (q_a, (a+2)^\wedge d\$x_i, 0^{n-i}\$v_i, z_i\$)$ occurs when state q_a is visited for the i th time. Since the run is single-pass on \mathbf{c} , we know that $v_i \leq_{\#} 0^i$ for all $i = 0, \dots, n$. Since it is single-pass on \mathbf{o} , we deduce that $x_0 \leq_{\#} y_d \dots y_{a+1}$, then $x_{i+1} \leq_{\#} x_i z_i$ for all i , and finally $x'_n \leq_{\#} x_n (a+2)^\wedge d$, with also $z_0 \leq_{\#} y_a (a+1)$. Finally, $z_{i+1} \leq_{\#} z_i$ since each subrun $C_i \xrightarrow{*}_w C_{i+1}$ is single-pass on \mathbf{t} .

All this yields $x'_n \leq_{\#} s(\lambda_n)$ and $v'_n \leq_{\#} 0^n$. Hence $S_{d,\lambda}$ is safe and robust: there is an execution $(p, s(\alpha)\$, 0^n\$, \$) \xrightarrow{*}_w (r, s(\beta), 0^{n'}\$, \$)$ if and only if $\alpha \in \text{Lim}$, $s(\beta) \leq_{\#} s(\alpha_n)$ and $n' \leq n$, entailing $H^\beta(n') \leq H^\alpha(n)$.

There remains to consider $S_{d,\lambda}^{-1}$, the PCS component that implements inverse Hardy steps for limits, see Figure 6. For given $a < d$, $S_{d,\lambda,a}^{-1}$ assumes that channel \mathbf{o} contains $s(\lambda_n) = y_d \dots y_{a+1} [y_a (a+1)]^n (a+2)^\wedge d$, guesses the position of the first $y_a (a+1)$ factor, and checks that it indeed occurs n times if \mathbf{c} contains 0^n . This check uses copies z_1, z_2, \dots of $y_a (a+1)$ temporarily stored on \mathbf{t} . Then $S_{d,\lambda}^{-1}$ writes back $s(\lambda) = y_d \dots y_{a+1} z a^\wedge d$ on \mathbf{o} , where $z(a+1) = z_n$. The reader should be easily convinced that, as far as one considers reliable steps, $S_{d,\lambda}^{-1}$ guarantees

$$(p, s(\alpha)\$, 0^n\$, \$) \xrightarrow{*}_{\text{rel}} (r, u, v, w) \text{ iff } \exists \lambda \in \text{Lim} : \alpha = \lambda_n, u = s(\lambda)\$, v = 0^n\$, \text{ and } w = \$. \quad (5.19)$$




 FIGURE 6. $S_{d,\lambda,a}^{-1}$, a PCS for inverse Hardy steps $(\lambda_n, n) \xrightarrow{H^{-1}} (\lambda, n)$.


FIGURE 7. Schematics for Theorem 5.8.

When superseding is taken into account, a run from p to r in $S_{d,\lambda}^{-1}$ has the form $(p, s(\alpha)\$, 0^n\$, \$) \xrightarrow{*}_w C_1 \xrightarrow{*}_w C_2 \xrightarrow{*}_w \dots C_n \xrightarrow{*}_w (r, u, v, w)$ where, for $i = 1, \dots, n$, C_i is the i th configuration that visits state q_a . Necessarily, C_i is some $(q_a, x_i\$x, 0^{n-i}0\$v_i, z_i\$)$. The first visit to q_a has $x \leq_{\#} y_d \dots y_{a+1}$, $z_1 \leq_{\#} y_a(a+1)$ and $v_1 = \varepsilon$, the following ones ensure $x_i = z_i x_{i+1}$, $z_{i+1} \leq_{\#} z_i$ and $v_{i+1} \leq_{\#} v_i 0$. Concluding the run requires $x_n = (a+2)^{\wedge} d$. Finally $v \leq_{\#} 0^n\$, s(\beta) = y_d \dots y_{a+1}(a+1)z_1 \dots z_{n-1}(a+2)^{\wedge} d$ and $u \leq_{\#} y_d \dots y_{a+1} z a^{\wedge} d$ for $z(a+1) = z_n \leq_{\#} z_{n-1} \leq_{\#} \dots z_2 \leq_{\#} z_1 \leq_{\#} y_a(a+1)$. Thus $u = s(\beta)\$$ and $v = 0^{n'}\$$ imply $s(\beta) \leq_{\#} s(\lambda)$ for some λ with $s(\lambda_n) \leq_{\#} s(\alpha)$, yielding $H^{\beta}(n') \leq H^{\lambda}(n) = H^{\lambda_n}(n) \leq H^{\alpha}(n)$.

5.4. Wrapping It Up. With the above weak Hardy computers, we have the essential gadgets required for our reductions. The wrapping-up is exactly as in [HSS12, Sch10b] (with a different encoding and a different machine model) and will only be sketched.

Theorem 5.8 (Verifying PCSs is Hard). *Reachability and Termination of PCSs are $\mathbf{F}_{\varepsilon_0}$ -hard.*

Proof. We exhibit a LOGSPACE reduction from the halting problem of a Turing machine M working in F_{ε_0} space to the Reachability problem in a PCS. We assume wlog. M to start in a state p_0 with an empty tape and to have a single halting state p_h that can only be reached after clearing the tape.

Figure 7 depicts schematically the PCS S we construct for the reduction. Let $n \stackrel{\text{def}}{=} |M|$ and $d \stackrel{\text{def}}{=} n + 1$. A run in S from the initial configuration to the final one goes through three stages:

- (1) The first stage robustly computes $F_{\varepsilon_0}(|M|) = H^{\Omega_d}(n)$ by first writing $s(\Omega_d)\$,$ i.e. $0^{\wedge}d\$,$ on \circ , $0^n\$,$ on \mathfrak{c} , and $\$$ on \mathfrak{t} , then by using S_d to perform forward Hardy steps; thus upon reaching state p_0 , \circ and \mathfrak{t} contain $\$$ and \mathfrak{c} encodes a budget $B \leq F_{\varepsilon_0}(|M|)$.
- (2) The central component simulates M over \mathfrak{c} where the symbols 0 act as blanks—this is easily done by cycling through the channel contents to simulate the moves of the head of M on its tape. Due to superseding steps, the outcome upon reaching p_h is that \mathfrak{c} contains $B' \leq B$ symbols 0 .
- (3) The last stage robustly computes $(F_{\varepsilon_0})^{-1}(B')$ by running S_d^{-1} to perform backward Hardy steps. This leads to \circ containing the encoding of some ordinal α and \mathfrak{c} of some n' , but we empty these channels and check that $\alpha = \Omega_d$ and $n' = n$ before entering state q_h .

Because $H^{\Omega_d}(n) \geq B \geq B' \geq H^{\alpha}(n') = H^{\Omega_d}(n)$, all the inequalities are actually equalities, and the simulation of M in stage 2 has necessarily employed reliable steps. Hence, M halts if and only if $(q_h, \varepsilon, \varepsilon, \varepsilon)$ is reachable from $(q_0, \varepsilon, \varepsilon, \varepsilon)$ in S .

The case of (non-)Termination is similar, but employs a *time* budget in a separate channel in addition to the space budget, in order to make sure that the simulation of M terminates in all cases, and leads to a state q_h that is the only one from which an infinite run can start in S . \square

6. ALTERNATIVE SEMANTICS

In this section we consider variant models for channel systems with priorities or losses and compare them with our PLCS model. The aim is to better understand the consequences, or lack thereof, of our choices.

We first consider *strict-superseding* systems, where messages may only supersede messages of strictly lower priority, and *overtaking* systems, where higher priority messages may move ahead of lower priority messages instead of erasing them. For completeness, we also discuss systems based on *priority queues*, where overtaking of lower priority messages is mandatory.

6.1. Strict-Superseding and Overtaking. In this section, we discuss two alternative operational semantics for PCSs that may seem more natural than our standard $\mathcal{S}_{\#}$.

Strict-Superseding Semantics: Here, a high-priority message may only supersede messages having *strictly* lower priority. Formally, we replace the internal-superseding relation

$C \xrightarrow{\#k}_{\#} C'$ with a new superseding relation, denoted $C \xrightarrow{\#k}_{\succ} C'$, and based on

$$x \xrightarrow{\#k}_{\succ} y \stackrel{\text{def}}{\iff} x = a_1 \cdots a_\ell \wedge y = a_1 \cdots a_{k-1} \cdot a_{k+1} \cdots a_\ell \wedge a_k < a_{k+1} .$$

Equivalently, we replace the rewrite rules from Eq. (2.1) with $\{a a' \rightarrow a' \mid 0 \leq a < a' \leq d\}$.

Overtaking semantics: Here, a high-priority message may move ahead of a low-priority message but this does not erase the low-priority message. Formally, we replace $C \xrightarrow{\#k}_{\#} C'$

with $C \xrightarrow{\#k}_{\text{ot}} C'$, based on

$$x \xrightarrow{\#k}_{\text{ot}} y \stackrel{\text{def}}{\iff} x = a_1 \cdots a_k a_{k+1} \cdots a_\ell \wedge y = a_1 \cdots a_{k+1} a_k \cdots a_\ell \wedge a_k < a_{k+1} .$$




FIGURE 8. Simulating reliable channels (left) with “strict-superseding” or “overtaking” PCSs (right).

In rewriting terms, \rightarrow_{ot} is defined by the rules $\{aa' \rightarrow a'a \mid 0 \leq a < a' \leq d\}$.

These two mechanisms drop fewer messages than our internal-superseding semantics. They may however be inadequate in case of network congestion, for instance they offer no solutions if all the messages in the congested buffers have the same priority. In any case, we show below that verification is undecidable for these two variant semantics (one can simulate Turing-powerful reliable channel systems with PCS under strict-superseding or overtaking semantics), which explains our choice of internal-superseding semantics.

Theorem 6.1. *Reachability and Termination are undecidable for PCSs under both the strict-superseding and overtaking semantics.*

Proof. We reduce from the reachability problem for channel systems with reliable channels which is known to be undecidable [BZ83]. A system S with reliable channels uses a finite (un-prioritized) alphabet $\Sigma = \{a_0, \dots, a_{p-1}\}$ and is equipped with m channels c_1, \dots, c_m . We simulate S with a PCS S' with strict superseding semantics having the same m channels and using the Σ_d priority alphabet with $d = p$. We use $d \in \Sigma_d$ as a separator, denoted $\$$ for clarity, while the other priorities $i \in \{0, \dots, p-1\}$ represent the original messages a_i . A string $w = a_{i_1} \dots a_{i_n} \in \Sigma^*$ is encoded as $\tilde{w} \stackrel{\text{def}}{=} i_1 \$ \dots i_n \$ \in \Sigma_d^*$. The actual reduction is obtained by equipping S' with transition rules that simulate the rules of S as illustrated in Figure 8. In essence, where S would write a_i , S' will write i followed by $\$$, *i.e.*, \tilde{a}_i , and S' will read $i' \cdot \$$ where S would read $a_{i'}$.

With the *strict superseding* policy, the only superseding that can occur is to have $\$$ erase a preceding $i < \$$. This results in a channel containing two or possibly more consecutive $\$$ symbols, a pattern that will never vanish in this simulation and that eventually forbids reading on the involved channel. In particular, any run of S' that reaches a final configuration $C_{\text{end}} = (q_{\text{end}}, \varepsilon, \dots, \varepsilon)$ has not used any strict superseding and thus corresponds to a run of the reliable channel system S . The same reduction works for Termination.

With the overtaking semantics, only $\$$ can overtake “original” messages of S in \tilde{w} . However, such an overtake results in having two consecutive $\$$ symbols on the channel, a pattern that can never disappear. Behind the $\$\$$ block, two lower messages $0 \leq i, j < d$ may occur consecutively and be open to overtaking but this will not derail the simulation since S' cannot read beyond $\$\$$. \square

6.2. Channels as Priority Queues. For the sake of completeness, let us mention channel systems where channels behave as priority queues. Here, reading from a channel will always read a message having highest priority among the contents of the channel. This can be seen as an extreme version of the overtaking semantics, where overtaking is mandatory. Such a model is not relevant for our purposes since it is not meant to handle congested

communication links: instead, it uses priorities as a way of choosing in which order messages should be processed.

Let us mention that finite-state communicating systems with priority queues can easily simulate Minsky machines by using a queue for each counter and two priorities: high-priority messages encode the counter value in unary, while a low-priority message can only be read in case of a zero-valued counter. They are hence Turing-powerful.

The most relevant case however is that of a single communication bus where many processes can read and write. Because only one queue is available—and still assuming a singleton alphabet for message contents—it is easy to see that priority queues systems are equivalent to Minsky counter machines restricted to *nested zero-tests*. Recall that for a machine with m counters c_1, \dots, c_m , nested zero-tests are tests of the form “ $(c_1 = 0 \wedge c_2 = 0 \wedge \dots \wedge c_i = 0)$?”, *i.e.*, one can only test the i th counter for emptiness when already the previous counters are empty. While Minsky counter machines with arbitrary zero-tests are Turing-powerful, they are equivalent to Petri nets when zero-tests are forbidden. Minsky machines with nested zero-tests are an intermediary model for which reachability is known to be decidable, see [Rei08] and [Bon13, Chapter 5].

7. HIGHER-ORDER LOSSY CHANNEL SYSTEMS

In this section, we introduce higher-order lossy channel systems, aka HOLCSs, a family of models that extend lossy channel systems. While a higher-order pushdown automaton has a stack of stacks of \dots of stacks [Aho69], a HOLCS has lossy channels inside lossy channels inside \dots inside lossy channels. (In this setting, the “dynamic” lossy channel systems from [AAC12] are a special case of 2LCSs, or second-order LCSs.) HOLCSs are well-structured, see Theorem 7.1, hence enjoy the usual decidability results from well-structured systems theory.

Our main result is that PCSs can simulate HOLCSs, see Theorem 7.7. On the one hand, this underlines the expressive power and naturalness of PCSs, in particular since the reductions we provide are quite straight-forward. On the other hand, we immediately obtain undecidability results: problems like boundedness or repeated control-state reachability are undecidable for PCSs since they already are for (first-order) LCSs (see [May03, Sch10a]).

7.1. Syntax. Formally, and given $k \in \mathbb{N}$, a k th-order LCS $S = (k, \Sigma, \text{Ch}, Q, \Delta)$ has first-order, second-order, \dots , up to k th-order, channels. We assume for simplicity that S has, for all $n = 1, \dots, k$, the same number m of n th-order channels, denoted $c_{n,1}, \dots, c_{n,m}$. Standardly, S uses a finite (un-prioritized) alphabet $\Sigma = \{a_1, \dots, a_p\}$. Write Σ^{*1} for the set of finite sequences of messages (usually written just Σ^*), and $\Sigma^{*(n+1)}$ for the set of finite sequences of elements from Σ^{*n} . We further order each $\Sigma^{*(n+1)}$ with $\leq_{*(n+1)}$, *i.e.*, the sequence extension of (Σ^{*n}, \leq_{*n}) , equating (Σ^{*0}, \leq_{*0}) with $(\Sigma, =)$. Precisely, given two sequences $x = x_1 \dots x_\ell$ and $y = y_1 \dots y_m$ in $\Sigma^{*(n+1)}$, we let

$$x_1 \dots x_\ell \leq_{*(n+1)} y_1 \dots y_m \stackrel{\text{def}}{\iff} \exists 1 \leq i_1 < i_2 < \dots < i_\ell \leq m : x_1 \leq_{*n} y_{i_1} \wedge \dots \wedge x_\ell \leq_{*n} y_{i_\ell}$$

Using Higman’s Lemma and induction over n , one sees that (Σ^{*n}, \leq_{*n}) is a well-quasi-order for any n .




FIGURE 9. Reading, writing, and losing in HOLCSs.

At any given time, the contents of a n th-order channel is a sequence $w \in \Sigma^{*n}$, so that a configuration of S has the form $C = (q, x_{1,1}, \dots, x_{1,m}, \dots, x_{k,1}, \dots, x_{k,m})$ with q a control state and $x_{n,i} \in \Sigma^{*n}$ for all $1 \leq n \leq k$ and $1 \leq i \leq m$. These configurations are ordered by

$$(q, x_{1,1}, \dots, x_{k,m}) \preceq_{\text{ho}} (q', y_{1,1}, \dots, y_{k,m}) \stackrel{\text{def}}{\iff} q = q' \wedge \forall n, i : x_{n,i} \leq_{*n} y_{n,i}.$$

This ordering of configurations is a wqo since, for each $n = 1, \dots, k$, (Σ^{*n}, \leq_{*n}) is.

7.2. Semantics. A HOLCS S as above has a *reliable* transition $C \xrightarrow{\delta} C'$ between configurations $C = (q, \dots, x_{i,n}, \dots)$ and $C' = (q', \dots, y_{i,n}, \dots)$ if $\delta = (q, \text{op}, q')$ is a rule moving control from q to q' and if the channel contents are modified according to the operation carried by δ , as we now define. There are four cases:

- op = $c_{1,i}!a$:** for some $a \in \Sigma$ and $1 \leq i \leq m$: then $y_{1,i} = x_{1,i} \cdot a$ while $y_{j,n} = x_{j,n}$ when $n > 1$ or $j \neq i$, *i.e.*, one writes a message to a 1st-order channel as in standard channel systems;
- op = $c_{1,i}?a$:** then one reads a message from a 1st-order channel, *i.e.*, $x_{1,i} = a \cdot y_{1,i}$ while the other channels are untouched;
- op = $c_{n+1,i}!c_{n,j}$:** for some $1 \leq n < k$: then one appends a copy of the whole contents of $c_{n,j}$ (a n th-order channel) to $c_{n+1,i}$, where it becomes the last element of the higher-order sequence. Formally, $y_{n+1,i} = x_{n+1,i} \cdot x_{n,j}$ and the other channels are untouched;
- op = $c_{n,i}?c_{n+1,j}$:** then one moves the first element of the higher-order sequence currently in $c_{n+1,j}$ to $c_{n,i}$ where it becomes the whole contents (the previous contents is erased). Formally, if $u \in \Sigma^{*n}$ is the first element of $x_{n+1,j}$, then $y_{n,i} = u$, $u \cdot y_{n+1,j} = x_{n+1,j}$, and the other channels are untouched.

In addition, all steps $C \rightarrow C'$ for $C' \preceq_{\text{ho}} C$, called *losing steps*, are allowed. This states that at any time the system may lose individual messages, sequences of messages, sequences of sequences of ... of messages, anywhere inside the channels.


Figure 9 illustrates the behavior of higher-order channels under reads, writes and losses (control states omitted).

Theorem 7.1. *HOLCSs equipped with \preceq_{ho} are well-structured.*

Proof. The ordering $C \preceq_{\text{ho}} C'$ entails $C' \xrightarrow{*} C$ (via losing steps). □

7.3. Simulation by PCSs. Let us consider a k th order LCS S with $\Sigma = \{a_1, \dots, a_p\}$ and $k \cdot m$ channels. We assume that $m = 1$ in order to simplify our constructions and proofs but they extend directly to the more expressive cases where $m > 1$.

We simulate S with a d -PCS \tilde{S} having k channels and $d \stackrel{\text{def}}{=} p + k - 1$. In Σ_d , the lower priorities $0, \dots, p - 1$ will be denoted a_1, \dots, a_p since they directly encode the messages from

FIGURE 10. Simulation of HOLCS S with PCS \tilde{S} : first-order rules.

Σ , while the higher priorities $p, \dots, p+k-1$ will be denoted $\$, \dots, \k since they are used as separators.

7.3.1. Encoding Configurations. A configuration $C = (q, x_1, \dots, x_k)$ of S is encoded as $\tilde{C} \stackrel{\text{def}}{=} (q, [x_1]_1, \dots, [x_k]_k)$, where, for $n = 1, \dots, k$, $[x]_n$ denotes the n th-level encoding of a n th-level sequence $x = u_1 u_2 \dots u_\ell \in \Sigma^{*n}$. Encodings are defined with

$$\begin{aligned} [a_{i_1} a_{i_2} \dots a_{i_\ell}]_1 &\stackrel{\text{def}}{=} \$1 a_{i_1} \$1 a_{i_2} \$1 \dots a_{i_\ell} \$1, \\ [u_1 u_2 \dots u_\ell]_{n+1} &\stackrel{\text{def}}{=} \$_{n+1} [u_1]_n \$_{n+1} [u_2]_n \$_{n+1} \dots [u_\ell]_n \$_{n+1}. \end{aligned}$$

Note that $[\varepsilon]_n = \$n$ and differs from $\varepsilon \in \Sigma_d^*$. For $n = 1, \dots, k$, we let $E_n \stackrel{\text{def}}{=} \{[x]_n \mid x \in \Sigma^{*n}\}$ denote the set of all n -level encodings. These are regular languages that are captured by the following regular expressions:


$$E_1 \stackrel{\text{def}}{=} \$1 \cdot ((a_1 + \dots + a_p) \cdot \$1)^*, \quad E_{n+1} \stackrel{\text{def}}{=} \$_{n+1} \cdot (E_n \cdot \$_{n+1})^*.$$

With this encoding, (Σ^{*n}, \leq_{*n}) and $(E_n, \leq_{\#})$ are isomorphic:

Lemma 7.2. *For any $x, y \in \Sigma^{*n}$, $x \leq_{*n} y$ if, and only if, $[x]_n \leq_{\#} [y]_n$.*

Proof Idea. By induction on n . For the “ \Leftarrow ” direction it is easier to rely on priority embedding, i.e., prove that $[x]_n \sqsubseteq_P [y]_n$ implies $x \leq_{*n} y$ and apply Lemma 3.3. \square

7.3.2. Encoding First-Order Rules. We may now complete the definition of \tilde{S} by describing its rules, with the goal of simulating the operational semantics of S while working on \tilde{C} and on encodings of sequences from some Σ^{*n} . This is easy for 1st-order rules that operate on c_1 only: where S writes a_i , \tilde{S} writes $a_i \cdot \$1$. Where it reads a_j , \tilde{S} reads $\$1 \cdot a_j$. This is illustrated in Figure 10.


 FIGURE 11. Simulation of HOLCS S with PCS \tilde{S} : higher-order rules.

7.3.3. Encoding Higher-Order Rules. Higher-order rules of the form $q_1 \xrightarrow{c_{n+1}!c_n} q_2$ are simulated in \tilde{S} as we illustrate in case δ_3 of Figure 11. Here \tilde{S} uses a loop, abbreviated as $\times \xrightarrow{c_n?u; c_n!u; c_{n+1}!u} \times$, to append a copy of c_n 's contents to c_{n+1} . This uses a high-priority $\$_{n+1}$ to mark the end of c_n 's contents and ensure that all of it has been read (and written back). Note that the loop checks that u is in E_n , *i.e.*, is a well-formed encoding, which is done by following a DFA for E_n . When the transfer is completed, a $\$_{n+1}$ must be appended to c_{n+1} to ensure consistency.

Simulating rules of the form $q_3 \xrightarrow{c_n?c_{n+1}} q_4$ follows the same logic (see Figure 11): \tilde{S} reads the first n -level encoding in c_{n+1} , checking that it is well-formed (with “ $u \in E_n$ ”) and writes it to c_n . Simultaneously, the previous contents of c_n is emptied by writing a $\$_{n+1}$ and reading it.

7.3.4. Correctness. The correctness of this simulation is captured by the next two propositions.

Proposition 7.3. *If S has a run $C \xrightarrow{*} C'$, then \tilde{S} has a run $\tilde{C} \xrightarrow{*} \tilde{C}'$.*

Proof Idea. On the one hand, \tilde{S} has been designed so that its behavior *without any superseding* directly mimics on encodings the effect of the reliable steps $C \rightarrow C'$. Then lossy steps in S are simulated by superseding since, thanks to Lemma 7.2, $C \succ_{\text{ho}} C'$ entails $\tilde{C} \geq_{\#} \tilde{C}'$. \square

There is an exact reciprocal to Proposition 7.3:

Proposition 7.4. *If \tilde{S} has a run $\tilde{C} \xrightarrow{*} \tilde{C}'$ then S has a run $C \xrightarrow{*} C'$.*

The correctness proof is harder in this direction since the steps of \tilde{S} are finer-grained than the steps of S . Note that a configuration $D = (q, w_1, \dots, w_k)$ of \tilde{S} is not necessarily the encoding \tilde{C} of a configuration of S : if q is not an original state of S (*i.e.*, is one of the unnamed states depicted on the right-hand side of Figure 10 or Figure 11) then D is not a \tilde{C} . Furthermore, if q is an original state, it is possible that some w_n does not belong to E_n .

With these difficulties in mind, we say that $D \in \text{Conf}_{\tilde{S}}$ is *safe* if every w_n ends with a $\$_n$ (and contains no $\$_{n'}$ for $n' > n$), and that D is a *stable* configuration if q is an original state.

We rely on the write-superseding semantics (see Section 2.1.3) for a better control on the form of the runs. We say that a write-superseding run $D_0 \rightarrow_w D_1 \rightarrow_w \dots \rightarrow_w D_r$ in \tilde{S} is a *macro-step* if D_0 and D_r are the only stable configurations it visits (in essence, a macro-step just follows the rules introduced in \tilde{S} to simulate a single rule of S using write-superseding semantics). We are now ready for the following lemmata.

Lemma 7.5. *Let $D \xrightarrow{\pm}_w D'$ be any macro-step in \tilde{S} . If D is safe then D' is safe too.*

Proof. Since the last message in a channel cannot be superseded, a c_n whose contents is safe remains safe if one does not read its final $\$n$, or one appends some safe contents. We now consider all four cases for the macro-step $D \xrightarrow{\pm}_w D'$.

- δ_1 : using rules of the form $q_1 \xrightarrow{c_1!a_i c_1!\$1} q_2$, it writes a safe $a_i \cdot \$1$ in c_1 (and does not read from the other channels).
- δ_2 : reading with some $q_3 \xrightarrow{c_1?\$1 c_1?a_j} q_4$, the read $\$1$ cannot be the final one in c_1 .
- δ_3 : using a macro-step that simulates $q_1 \xrightarrow{c_{n+1}!c_n} q_2$, the last write to c_n is a safe $u \in E_n$, and the last write to c_{n+1} is $\$(n+1)$.
- δ_4 : using a macro-step that simulates $q_3 \xrightarrow{c_n?c_{n+1}} q_4$, the last write to c_n is a safe $u \in E_n$, which is also the last read from c_{n+1} , implying that the final $\$(n+1)$ in c_{n+1} cannot have been read. \square

Lemma 7.6. *If D is safe and $D \xrightarrow{\pm}_w \tilde{C}'$ is a write-superseding macro-step in \tilde{S} , then there is a $C \in \text{Conf}_S$ such that $D \geq_{\#} \tilde{C}$ and $C \xrightarrow{\pm} C'$ in S .*

Proof. Write $D = (q, w_1, \dots, w_k)$ and $\tilde{C}' = (q', v_1, \dots, v_k)$. For each $n = 1, \dots, k$, we know that w_n is some $w'_n \$n$ (since D is safe) and that v_n is some $\lfloor y_n \rfloor_n$. We now consider four cases for the macro-step $D \xrightarrow{\pm}_w \tilde{C}$:

- δ_1 : it uses some $q = q_1 \xrightarrow{c_1!a_i c_1!\$1} q_2 = q'$. From the definition of \rightarrow_w (see Section 2.1.3) we deduce that $v_1 = x \cdot \$1$ where x is a prefix of $w'_1 \cdot \$1 \cdot a_i$ while $w_n = v_n$ for $n > 1$. Since $v_1 \in E_1$, either $v_1 = w_1 \cdot a_i \cdot \1 and we take $\tilde{C} = D$, or v_1 is a safe prefix of w_1 , in which case we take $C = (q, y_1, \dots, y_n)$.
- δ_2 : it uses some $q = q_3 \xrightarrow{c_1?\$1 c_1?a_j} q_4 = q'$ with no writing. The write-superseding semantics entails $w_1 = \$1 \cdot a_j \cdot v_1$. Here $w_1 = \lfloor a_j \cdot y_1 \rfloor_1$ and we take $\tilde{C} = D$.
- δ_3 : the macro-step writes a $\$(n+1)$ to c_n and reads it back, so that $u = w_n$ and we deduce that $w_n \in E_n$ and is some $\lfloor x \rfloor_n$. Furthermore v_n is u perhaps after some superseding and one obtains $x \geq_{*n} y_n$ from $u \geq_{\#} v_n$.

On c_{n+1} the macro-step writes $u \cdot \$(n+1)$ and we reason as in case δ_1 : if write-superseding erases the $\$(n+1)$ that closes w_{n+1} then $v_{n+1} \leq_{\#} w_{n+1}$ and we set $C = (q, \dots, y_{n-1}, x, y_{n+1}, \dots)$, otherwise $v_{n+1} = w_{n+1} \cdot \lfloor x' \rfloor_n \$(n+1)$ with $\lfloor x' \rfloor_n \leq_{\#} u$, we know that w_{n+1} is some $\lfloor x_{n+1} \rfloor_{n+1}$ and we may set $C = (q, \dots, y_{n-1}, x, x_{n+1}, \dots)$.

- δ_4 : On c_{n+1} the macro-step just reads $\$(n+1) \cdot u$, where $u \in E_n$ is some $\lfloor x \rfloor_n$. Necessarily w_n is $\lfloor x \cdot y_{n+1} \rfloor_{n+1}$. On c_n , one writes $\$(n+1) \cdot u$ and reads $\$(n+1)$. Necessarily $y_n \leq_{\#} u$, hence $y_n \preceq_{\text{ho}} x$. Taking $C = (q, \dots, y_{n-1}, \varepsilon, x \cdot y_{n+1}, y_{n+2}, \dots)$ works. \square

We can now conclude our correctness proof.

Proof of Proposition 7.4. Assume $\tilde{C} \xrightarrow{*} \tilde{C}'$. We first apply Lemma 2.8 and deduce the existence of a write-superseding run $D_0 \xrightarrow{\pm}_w \tilde{C}'$ for some $D_0 \leq_{\#} \tilde{C}$. Let us single out the

stable configurations along this run and write it under the form $D_0 \xrightarrow{\pm}_w D_1 \xrightarrow{\pm}_w \cdots D_{r-1} \xrightarrow{\pm}_w D_r = \widetilde{C}'$, *i.e.*, as a sequence of r macro-steps.

We now reason by induction on r . If $r = 0$, then $D_0 = \widetilde{C}'$ and $\widetilde{C} \geq_{\#} \widetilde{C}'$, implying $C \succ_{\text{ho}} C'$ by Lemma 7.2 so that S has $C \xrightarrow{*} C'$ via lossy steps.

If $r > 0$, we first observe that D_0 is safe (since \widetilde{C} is) hence D_1, \dots, D_r too by Lemma 7.5. Using Lemma 7.6 on $D_{r-1} \xrightarrow{\pm}_w \widetilde{C}'$ implies that there is some $C'' \in \text{Conf}_S$ with $D_{r-1} \geq_{\#} \widetilde{C}''$ and $C'' \xrightarrow{\pm} C'$. On the other hand, the run $\widetilde{C} \xrightarrow{*} D_0 \xrightarrow{\pm}_w D_{r-1} \xrightarrow{*} \widetilde{C}''$ can be transformed into $r - 1$ macro-steps. We can thus apply the ind. hyp. and deduce that $C \xrightarrow{*} C''$ in S , \square

We may now state formally the main theorem of this section.

Theorem 7.7. *There is a LOGSPACE reduction that transforms reachability problems on k th-order HOLCSs to reachability problems on PCSs of level $d = k + 1$.*

Proof Sketch. The above ideas use $p + k$ priority levels. One can tighten these simulations to use fewer priorities and complete the proof of Theorem 7.7 by encoding the messages a_0, \dots, a_{p-1} as *fixed length* binary strings over $\{0, 1\}$ followed by a $\$$ separator. Then the prioritized alphabet $\{0, 1, \$, \dots, \$\}$ with $k + 2$ priority levels suffices.

In particular $\{0, 1, \$\}$ is enough for LCSs. In the case of *weak* LCSs where the set of messages is linearly ordered (say $a_0 < a_1 < \dots < a_{p-1}$) and where, in addition to message losses, any message can be replaced by a lower message inside the channels, we can further tighten this to $\{0, \$\}$ with a unary encoding of message a_i as $0^i \$$. \square

Remark 7.8. The simulation of HOLCSs by PCSs is quite straightforward. We believe that a reduction from PCS reachability to HOLCS reachability must exist (on complexity-theoretical grounds) but we do not have at the moment any suggestion for a simple encoding of PCSs in HOLCSs. \square

8. APPLICATIONS OF THE PRIORITY EMBEDDING TO TREES

In this section we show how tree orderings can be reflected into sequences over a priority alphabet. This illustrates the “power” of priority embeddings, and yields as a byproduct a proof that strong tree embeddings form a wqo. The consequence will be that PCSs can perform operations on encodings of trees in a *robust* way, *i.e.*, such that superseding steps respect the tree ordering. This was already the key insight

- in Section 5 when we encoded ordinals, seen as terms, into sequences over Σ_d^* , and
- in Section 7 when we encoded nested sequences over Σ^{n*} , which can be seen as terms of bounded depth, also into sequences over Σ_d^* .

We generalize these ideas here, and allow in particular for labeled trees. More precisely, we show

- (1) in Section 8.1 that priority embeddings reflect the so-called *strong tree-embedding* ordering on trees of bounded depth, and
- (2) in Section 8.2 that they reflect the *immersion ordering* used by Gupta [Gup92] to compute the maximal order type of the *tree minor ordering*.

These two tree orderings are already known to be wqos—with a maximal order type ε_0 that matches that of priority embeddings. The point of the section is thus rather to show how to encode trees robustly in priority strings, and thus how to manipulate trees robustly in PCSs.

Remark 8.1 (Kruskal’s Tree Theorem). The reader is likely to be already acquainted with the homeomorphic embedding ordering used in Kruskal’s Tree Theorem. The strong tree embedding refines this ordering but is restricted to trees of bounded depth (it is not a wqo on general trees), while the tree minor ordering is coarser than the homeomorphic embedding ordering. Here, we wish to explicitly emphasize that reflections in priority embeddings do *not* provide a new proof of Kruskal’s Tree Theorem. In fact, such a proof cannot exist since the maximal order type for homeomorphic embeddings is considerably larger than ε_0 [RW93].

8.1. Reflecting Bounded Depth Trees. Recall Fact 3.5 about order reflections: if B reflects A and (B, \leq_B) is a wqo, then (A, \leq_A) is necessarily a wqo. Our goal is to show that $(\Sigma_{d,\Gamma}, \sqsubseteq_{p,\Gamma})$ reflects Γ -labeled trees of depth at most $d + 1$ endowed with the *strong tree-embedding* relation.

Given an alphabet Γ , we note $T(\Gamma)$ for the set of finite, ordered, unranked labeled trees (aka variadic terms) over Γ . Let d be a depth in \mathbb{N} . We work here with the set $T_d(\Gamma)$ of *trees of depth at most d* over Γ . Formally, $T_d(\Gamma)$ is defined by induction over d by $T_0(\Gamma) \stackrel{\text{def}}{=} \emptyset$, and for $d > 0$, $T_d(\Gamma)$ is the smallest set containing $T_{d-1}(\Gamma)$ and such that, if t_1, \dots, t_n are trees in $T_{d-1}(\Gamma)$ and f is in Γ , then the tree $f(t_1 \cdots t_n)$ obtained by adding an f -labeled root over them is in $T_d(\Gamma)$. When $n = 0$ we write f rather than $f()$. Figure 12 presents two labeled trees of depth 3.




FIGURE 12. Two trees in $T_3(\{a, b, f, g, r\})$.

It will be convenient in the following to use the *extension* operation “@” on trees, which is defined for $n \geq 0$ by

$$f(t_1 \cdots t_n) @ t \stackrel{\text{def}}{=} f(t_1 \cdots t_n t); \quad (8.1)$$

in particular, $f @ t = f(t)$. For instance, t_1 in Figure 12 can be decomposed as $r @ (f @ a) @ (g @ b)$.

In case where Γ is a singleton, we denote by “ \bullet ” its only element and write T_d for $T_d(\{\bullet\})$. For instance, $T_1 = \{\bullet\}$ contains a single tree.

8.1.1. Strong Tree Embeddings. Assume that (Γ, \leq_Γ) is a wqo, and that we have already defined a well-quasi-ordering \sqsubseteq_T on trees of maximal depth d —note that as a base case, since $T_0(\Gamma)$ is empty, it is vacuously well-quasi-ordered by the empty relation. We can lift it into a wqo $(T_{d+1}(\Gamma), \sqsubseteq_T)$ on trees of maximal depth $d + 1$ by $f(t_1 \cdots t_n) \sqsubseteq_T f'(t'_1 \cdots t'_m) \stackrel{\text{def}}{\Leftrightarrow} f \leq_\Gamma f'$ and $t_1 \cdots t_n \sqsubseteq_{T^*} t'_1 \cdots t'_m$, *i.e.*, by considering a product between (Γ, \leq_Γ) and the sequence embedding ordering on tree sequences $(T_d(\Gamma)^*, \sqsubseteq_{T^*})$: by Dickson’s Lemma and Higman’s Lemma, this defines a wqo on trees for every finite d , which we call the *strong tree embedding*. Put differently $t \sqsubseteq_T t'$ if t it can be obtained from t' by deleting whole subtrees and/or decrementing node labels.

Strong tree embeddings refine the *homeomorphic tree embeddings* used in Kruskal's Tree Theorem; in general they do not give rise to a wqo, but in the case of bounded depth trees they do. The two trees t_1 and t_2 in Figure 12 are *not* related by any homeomorphic tree embedding, and thus neither by strong tree embedding. The tree t_2 homeomorphically embeds into $t_3 = r(g(g(a)))$, but does not strongly embed into t_3 .

Observe that \sqsubseteq_T is a precongruence for $@$:

$$t_1 \sqsubseteq_T t'_1 \text{ and } t_2 \sqsubseteq_T t'_2 \text{ imply } t_1 @ t_2 \sqsubseteq_T t'_1 @ t'_2, \quad (8.2)$$

$$t \sqsubseteq_T t @ t'. \quad (8.3)$$

8.1.2. *Encoding Trees as Strings.* It is easy to encode trees into finite sequences. For instance, drawing inspiration from the ordinal encodings employed in Section 5, one might be tempted to encode the two trees in Figure 12 by

$$s_2(t_1) = (0, a)(1, f)(0, b)(1, g)(2, r) \quad (8.4)$$

$$s_2(t_2) = (0, a)(1, g)(2, r) \quad (8.5)$$

as the result of an inductive encoding $s_d(w(t_1 \cdots t_n)) \stackrel{\text{def}}{=} s_{d-1}(t_1) \cdots s_{d-1}(t_n) \cdot (d, w)$. Observe however that we wish our encoding to be an order reflection (*cf.* Section 3.2), which is not the case with s_d : we see that $s_2(t_2) \sqsubseteq_{p,\Gamma} s_2(t_1)$, although $t_2 \not\sqsubseteq_T t_1$. Over a singleton alphabet however, s_d is an order reflection from T_{d+1} to Σ_d^* .

Here we present a more redundant encoding, apt to handle arbitrary alphabets. We encode trees of bounded depth using the mapping $[\cdot]_d: T_{d+1}(\Gamma) \rightarrow \Sigma_{d,\Gamma}^*$ defined by induction on d by

$$[w(t_1 \cdots t_n)]_d \stackrel{\text{def}}{=} \begin{cases} (d, w) & \text{if } n = 0, \\ [t_1]_{d-1} \cdot (d, w) \cdots [t_n]_{d-1} \cdot (d, w) & \text{otherwise.} \end{cases} \quad (8.6)$$

For instance, if we fix $d = 2$, the trees in Figure 12 are encoded as

$$[t_1]_2 = (0, a)(1, f)(2, r)(0, b)(1, g)(2, r), \quad (8.7)$$

$$[t_2]_2 = (0, a)(1, g)(2, r). \quad (8.8)$$

This satisfies $[t_2]_2 \not\sqsubseteq_{p,\Gamma} [t_1]_2$ as desired.

8.1.3. *Proper Words.* Not every string in $\Sigma_{d,\Gamma}^*$ is the encoding of a tree according to $[\cdot]_d$: for $0 \leq a \leq d$, we let

$$P_{-1,\Gamma} \stackrel{\text{def}}{=} \emptyset, \quad P_{a,\Gamma} \stackrel{\text{def}}{=} \bigcup_{w \in \Gamma} (P_{a-1,\Gamma} \cdot \{(a, w)\})^* \cdot (P_{a-1,\Gamma} \cup \{\varepsilon\}) \cdot \{(a, w)\} \quad (8.9)$$

be the set of *proper encodings of height a*. Then $P \stackrel{\text{def}}{=} \bigcup_{a \leq d} P_{a,\Gamma}$ is the set of *proper words* in $\Sigma_{d,\Gamma}^*$. A proper word x belongs to a unique $P_{a,\Gamma}$ with $a = h(x)$, where $h(x)$ is the height of x , and has then a canonical factorization of the form $x = x_1(a, w) \cdots x_m(a, w)$ with every x_j in $P_{a-1,\Gamma}$ and w in Γ .

Given a depth d , we see that $[\cdot]_d$ is a bijection between $T_{d+1}(\Gamma)$ and $P_{d,\Gamma}$, with inverse $\tau: P_{d,\Gamma} \rightarrow T_{d+1}(\Gamma)^*$ defined using canonical decompositions by

$$\tau(x = x_1(h(x), w) \cdots x_m(h(x), w)) \stackrel{\text{def}}{=} w(\tau(x_1) \cdots \tau(x_m)) \quad (8.10)$$

$$= w @ \tau(x_1) @ \cdots @ \tau(x_m). \quad (8.11)$$

Proposition 8.2. *The map $[\cdot]_d$ is an order reflection from $(T_{d+1}^*, \sqsubseteq_{T^*})$ to $(\Sigma_d^*, \sqsubseteq_p)$.*

Proof. Let x and x' be two proper words in $P_{d,\Gamma}$ with $x \sqsubseteq_{p,\Gamma} x'$; we show by induction on x that $\tau(x) \sqsubseteq_T \tau(x')$. We consider the canonical factorizations $x = x_1(d, w) \cdots x_m(d, w)$ and $x' = x'_1(d, w') \cdots x'_n(d, w')$ for $m, n \geq 0$, x_j, x'_j in $P_{d-1,\Gamma}$ for all j , and w, w' in Γ .

By definition of the generalized priority embedding, the m pairs (d, w) occurring in x must be mapped to some pairs (d, w') occurring in x' with $w \leq_\Gamma w'$. Hence there exist $1 \leq i_1, \dots, i_m \leq n$ such that $x_j \sqsubseteq_{p,\Gamma} x'_{i_j}$. Therefore,

$$\begin{aligned} \tau(x) &= w @ \tau(x_1) @ \cdots @ \tau(x_m) && \text{by (8.11)} \\ &\sqsubseteq_T w' @ \tau(x_1) @ \cdots @ \tau(x_m) && \text{by (8.2) since } w \leq_\Gamma w' \\ &\sqsubseteq_T w' @ \tau(x'_{i_1}) @ \cdots @ \tau(x'_{i_m}) && \text{by (8.2) and ind. hyp. on } x_j \sqsubseteq_{p,\Gamma} x'_{i_j} \\ &\sqsubseteq_T w' @ \tau(x'_1) @ \cdots @ \tau(x'_n) && \text{by (8.3)} \\ &= \tau(x') && \text{by (8.11).} \quad \square \end{aligned}$$

Note that Proposition 8.2 provides an alternative proof of the fact that $(T_d(\Gamma), \sqsubseteq_T)$ is a wqo, thanks to Theorem 3.6 and Fact 3.5.

8.2. Relationship to Tree Minors. As a further application of the priority embedding, we demonstrate that we also subsume another wqo on trees, the *tree minor* ordering, by using the techniques of [Gup92] to encode trees into generalized prioritized alphabets. The tree minor ordering is coarser than the homeomorphic embedding, but the upside is that trees of *unbounded* depth can be encoded into strings.

The trees considered in [Gup92] are unlabeled finite rooted trees with an ordering on the children of every internal vertex, called *planar planted trees* therein, *i.e.* trees from $T \stackrel{\text{def}}{=} T(\{\bullet\})$. Figure 13 illustrates two such trees. The ordering on the children in particular implies that, for instance, the tree $\bullet(\bullet(\bullet), \bullet)$ is not equivalent to the tree $\bullet(\bullet, \bullet(\bullet))$. Gupta gives in [Gup92] a constructive proof that planar planted trees are well-quasi-ordered under minors. Recall that t_1 is a *minor* of t_2 if t_1 can be obtained from t_2 by a series of edge contractions, *e.g.* in Figure 13, the left-hand tree is a minor of the right-hand one. Note that, however, the two trees are incomparable for the previously considered homeomorphic embeddings.



FIGURE 13. Two trees in T_2 .

Gupta provides in [Gup92] an effective linearization $\text{lin}: T \rightarrow \bigcup_{d \geq 0} \Sigma_{d,\Gamma}^*$ which associates with every tree t a word $\text{lin}(t)$ over the generalized prioritized alphabet $\Sigma_{d,\Gamma}$, where d is

dubbed the *width* of t —which is at most its number of vertices—and $(\Gamma, =)$ is a finite alphabet with $\Gamma = \{0, 1, 2, 3\}$. Gupta continues by defining a so-called *immersion ordering* \sqsubseteq_I on $\Sigma_{d,\Gamma}^*$ for any fixed $d \in \mathbb{N}$ as follows: given $x = (a_1, w_1)(a_2, w_2) \cdots (a_k, w_k) \in \Sigma_{d,\Gamma}^*$ and $y \in \Sigma_{d,\Gamma}^*$, $x \sqsubseteq_I y$ if y can be factored as $y = y_0 y_1 \cdots y_k y_{k+1}$ such that

$$y_i \in (\Sigma_{d,\Gamma} \setminus \Sigma_{d-a_i-1,\Gamma})^* \cdot (a_i, w_i) \cdot (\Sigma_{d,\Gamma} \setminus \Sigma_{d-a_i-1,\Gamma})^*, \quad 1 \leq i \leq k.$$

The crucial relationship between trees in T , lin and \sqsubseteq_I is established in [Gup92, Theorem 4.1]: given planar planted trees $t_1, t_2 \in T$, whenever $\text{lin}(t_1) \sqsubseteq_I \text{lin}(t_2)$ then t_1 is a minor of t_2 . By showing that \sqsubseteq_I is a well-quasi-ordering, Gupta concludes that planar planted trees are well-quasi-ordered under minors.

The immersion ordering \sqsubseteq_I is closely related to our generalized priority ordering. In fact, it is easily seen that $\sqsubseteq_{p,\Gamma}$ can be viewed as a sub-structure of \sqsubseteq_I . Define an automorphism $\kappa : \Sigma_{d,\Gamma} \rightarrow \Sigma_{d,\Gamma}$ as

$$\kappa(a, w) \stackrel{\text{def}}{=} (d - a, w)$$

which canonically extends to words over $\Sigma_{d,\Gamma}^*$. Now $x \sqsubseteq_{p,\Gamma} y$ in particular implies $\kappa(x) \sqsubseteq_I \kappa(y)$. Thus, our Theorem 3.6 yields as a corollary another proof that the immersion ordering \sqsubseteq_I is a wqo.

Corollary 8.3. *The immersion ordering \sqsubseteq_I is a well-quasi-ordering.*

8.3. Further Applications. As stated in the introduction to this section, our main interest in strong tree embeddings is in connection with structural orderings of ordinals; see Section 5. Bounded depth trees are also used in the verification of infinite-state systems as a means to obtain decidability results, in particular for tree pattern rewriting systems [GMSZ08] in XML processing, and, using elimination trees [OdMN12], for bounded-depth graphs used *e.g.* in the verification of *ad-hoc* networks [DSZ10], the π -calculus [Mey08], programs [BKWZ13], and protocols [KS14]. These applications consider *labeled* trees, which motivate the generalized priority alphabets and embedding defined in Section 3.

The exact complexity of verification problems in the aforementioned models is currently unknown [GMSZ08, DSZ10, Mey08, BKWZ13, KS14]. Our encoding suggests they might be $\mathbf{F}_{\varepsilon_0}$ -complete. We hope to see PCS Reachability employed as a “master” problem for $\mathbf{F}_{\varepsilon_0}$, like LCS Reachability for $\mathbf{F}_{\omega^\omega}$, which is used in reductions instead of more difficult proofs based on Turing machines and Hardy computations.

9. CONCLUDING REMARKS

We introduced Priority Channel Systems, a natural model for protocols and programs with differentiated, prioritized asynchronous communications, and showed how they give rise to well-structured systems with decidable model-checking problems.

We showed that Reachability and Termination for PCSs are $\mathbf{F}_{\varepsilon_0}$ -complete, and we expect our techniques to be transferable to other models, *e.g.* models based on wqos on bounded-depth trees or graphs, whose complexity has not been analyzed [GMSZ08, DSZ10, Mey08, BKWZ13, KS14]. This is part of our current research agenda on complexity for well-structured systems [SS12].

In spite of their enormous worst-case complexity, we expect PCSs to be amenable to regular model checking techniques *à la* [AJ96, BG99]. This requires investigating the

algorithmics of upward- and downward-closed sets of configurations wrt. the priority ordering. These sets, which are always regular, seem promising since \sqsubseteq_p shares some good properties with the better-known subword ordering, *e.g.* the upward- or downward-closure of a sequence $x \in \Sigma_d^*$ can be represented by a deterministic finite automaton with $|x|$ states.

ACKNOWLEDGMENTS

We thank Lev Beklemishev who drew our attention to [SS85] and the reviewers for their helpful suggestions.

REFERENCES

- [AAC12] P. A. Abdulla, M. F. Atig, and J. Cederberg. Timed lossy channel systems. In *FST&TCS 2012*, volume 18 of *Leibniz International Proceedings in Informatics*, pages 374–386. Leibniz-Zentrum für Informatik, 2012.
- [AČJT00] P. A. Abdulla, K. Čerāns, B. Jonsson, and Yih-Kuen Tsay. Algorithmic analysis of programs with well quasi-ordered domains. *Information and Computation*, 160(1–2):109–127, 2000.
- [ADOW05] P. A. Abdulla, J. Deneux, J. Ouaknine, and J. Worrell. Decidability and complexity results for timed automata via channel machines. In *ICALP 2005*, volume 3580 of *Lecture Notes in Computer Science*, pages 1089–1101. Springer, 2005.
- [Aho69] A. V. Aho. Nested stack automata. *Journal of the ACM*, 16(3):383–406, 1969.
- [AJ96] P. A. Abdulla and B. Jonsson. Verifying programs with unreliable channels. *Information and Computation*, 127(2):91–101, 1996.
- [BBC⁺98] S. Blake, D. Black, M. Carlson, E. Davies, Z. Wang, and W. Weiss. An architecture for differentiated services. RFC 2475, December 1998.
- [BG99] B. Boigelot and P. Godefroid. Symbolic verification of communication protocols with infinite state spaces using QDDs. *Formal Methods in System Design*, 14(3):237–255, 1999.
- [BH99] A. Bouajjani and P. Habermehl. Symbolic reachability analysis of FIFO-channel systems with nonregular sets of configurations. *Theoretical Computer Science*, 221(1–2):211–250, 1999.
- [BJW82] R. V. Book, M. Jantzen, and C. Wrathall. Monadic Thue systems. *Theoretical Computer Science*, 19:231–251, 1982.
- [BKWZ13] K. Bansal, E. Koskinen, T. Wies, and D. Zufferey. Structural counter abstraction. In *TACAS 2013*, volume 7795 of *Lecture Notes in Computer Science*, pages 62–77. Springer, 2013.
- [BMO⁺12] P. Bouyer, N. Markey, J. Ouaknine, Ph. Schnoebelen, and J. Worrell. On termination and invariance for faulty channel machines. *Formal Aspects of Computing*, 24(4–6):595–607, 2012.
- [Bon13] R. Bonnet. *Theory of Well-Structured Transition Systems and Extended Vector-Addition Systems*. Thèse de doctorat, ENS Cachan, France, 2013.
- [BZ83] D. Brand and P. Zafropulo. On communicating finite-state machines. *Journal of the ACM*, 30(2):323–342, 1983.
- [CF05] G. Cécé and A. Finkel. Verification of programs with half-duplex communication. *Information and Computation*, 202(2):166–190, 2005.
- [CFP96] G. Cécé, A. Finkel, and S. Purushothaman Iyer. Unreliable channels are easier to verify than perfect channels. *Information and Computation*, 124(1):20–31, 1996.
- [CS08] P. Chambart and Ph. Schnoebelen. The ordinal recursive complexity of lossy channel systems. In *LICS 2008*, pages 205–216. IEEE Press, 2008.
- [dJP77] D. H. J. de Jongh and R. Parikh. Well-partial orderings and hierarchies. *Indagationes Mathematicae*, 39(3):195–207, 1977.
- [DSZ10] G. Delzanno, A. Sangnier, and G. Zavattaro. Parameterized verification of ad hoc networks. In *Concur 2010*, volume 6269 of *Lecture Notes in Computer Science*, pages 313–327. Springer, 2010.
- [FS01] A. Finkel and Ph. Schnoebelen. Well-structured transition systems everywhere! *Theoretical Computer Science*, 256(1–2):63–92, 2001.
- [FW98] M. Fairtlough and S. S. Wainer. Hierarchies of provably recursive functions. In S. Buss, editor, *Handbook of Proof Theory*, chapter III, pages 149–207. Elsevier, 1998.

- [GMSZ08] B. Genest, A. Muscholl, O. Serre, and M. Zeitoun. Tree pattern rewriting systems. In *ATVA 2008*, volume 5311 of *Lecture Notes in Computer Science*, pages 332–346. Springer, 2008.
- [Gup92] A. Gupta. A constructive proof that trees are well-quasi-ordered under minors. In *LFCS 1992*, volume 620 of *Lecture Notes in Computer Science*, pages 174–185. Springer, 1992.
- [HSS12] S. Haddad, S. Schmitz, and Ph. Schnoebelen. The ordinal-recursive complexity of timed-arc Petri nets, data nets, and other enriched nets. In *LICS 2012*, pages 355–364. IEEE Press, 2012.
- [HSS13] C. Haase, S. Schmitz, and Ph. Schnoebelen. The power of priority channel systems. In *Concur 2013*, volume 8052 of *Lecture Notes in Computer Science*, pages 319–333. Springer, 2013.
- [KS14] B. König and J. Stückrath. A general framework for well-structured graph transformation systems. In *Concur 2014*, volume 8704 of *Lecture Notes in Computer Science*, pages 467–481. Springer, 2014.
- [Kur06] A. Kurucz. Combining modal logics. In *Handbook of Modal Logics*, chapter 15, pages 869–926. Elsevier, 2006.
- [LB92] J.-Y. Le Boudec. The Asynchronous Transfer Mode: a tutorial. *Computer Networks and ISDN Systems*, 24(4):279–309, 1992.
- [LW70] M.H. Löb and S.S. Wainer. Hierarchies of number theoretic functions, I. *Archiv für Mathematische Logik und Grundlagenforschung*, 13:39–51, 1970.
- [LW08] S. Lasota and I. Walukiewicz. Alternating timed automata. *ACM Transactions on Computational Logic*, 9(2), 2008.
- [May03] R. Mayr. Undecidable problems in unreliable computations. *Theoretical Computer Science*, 297(1–3):337–354, 2003.
- [Mey08] R. Meyer. On boundedness in depth in the π -calculus. In *IFIP TCS 2008*, volume 273 of *IFIP*, pages 477–489. Springer, 2008.
- [Mus10] A. Muscholl. Analysis of communicating automata. In *LATA 2010*, volume 6031 of *Lecture Notes in Computer Science*, pages 50–57. Springer, 2010.
- [OdMN12] P. Ossona de Mendez and J. Nešetřil. *Sparsity*, chapter 6: Bounded height trees and tree-depth, pages 115–144. Springer, 2012.
- [OW07] J. Ouaknine and J. Worrell. On the decidability and complexity of Metric Temporal Logic over finite words. *Logical Methods in Computer Science*, 3(1):1–27, 2007.
- [Pac87] J. K. Pachl. Protocol description and analysis based on a state transition model with channel expressions. In *PSTV '87*, pages 207–219. North-Holland, 1987.
- [Rei08] K. Reinhardt. Reachability in Petri nets with inhibitor arcs. In *RP 2008*, volume 223 of *Electronic Notes in Theoretical Computer Science*, pages 239–264. Elsevier, 2008.
- [RW93] M. Rathjen and A. Weiermann. Proof-theoretic investigations on Kruskal’s Theorem. *Annals of Pure and Applied Logic*, 60(1):49–88, 1993.
- [Sch10a] Ph. Schnoebelen. Lossy counter machines decidability cheat sheet. In *RP 2010*, volume 6227 of *Lecture Notes in Computer Science*, pages 51–75. Springer, 2010.
- [Sch10b] Ph. Schnoebelen. Revisiting Ackermann-hardness for lossy counter machines and reset Petri nets. In *MFCS 2010*, volume 6281 of *Lecture Notes in Computer Science*, pages 616–628. Springer, 2010.
- [Sch13] S. Schmitz. Complexity hierarchies beyond Elementary. Preprint, December 2013. arXiv:1312.5686 [cs.CC].
- [SS85] K. Schütte and S. G. Simpson. Ein in der reinen Zahlentheorie unbeweisbarer Satz über endliche Folgen von natürlichen Zahlen. *Archiv für Mathematische Logik und Grundlagenforschung*, 25(1):75–89, 1985.
- [SS11] S. Schmitz and Ph. Schnoebelen. Multiply-recursive upper bounds with Higman’s Lemma. In *ICALP 2011*, volume 6756 of *Lecture Notes in Computer Science*, pages 441–452. Springer, 2011.
- [SS12] S. Schmitz and Ph. Schnoebelen. Algorithmic aspects of WQO theory. Lecture notes, 2012. cel.archives-ouvertes.fr:cel-00727025.
- [SS13] S. Schmitz and Ph. Schnoebelen. The power of well-structured systems. In *Concur 2013*, volume 8052 of *Lecture Notes in Computer Science*, pages 5–24. Springer, 2013. Invited talk.